FARMERS' BEHAVIOUR IN THE FACE OF UNCERTAINTY:
A BEHAVIOURAL ECONOMICS APPROACH TO FARMERS' LAND AND INSURANCE DECISIONS

By

Mary Kathryn Doidge

A DISSERTATAION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Agricultural, Food and Resource Economics – Doctor of Philosophy

2019

ABSTRACT

FARMERS' BEHAVIOUR IN THE FACE OF UNCERTAINTY:
A BEHAVIOURAL ECONOMICS APPROACH TO FARMERS' LAND AND INSURANCE DECISIONS

By

Mary Kathryn Doidge

Neoclassical economic theory hinges on the assumption of rational behaviour. The study of agents' behaviour in the face of risk and uncertainty is often guided by this assumption: agents are assumed to make predictions based on the information available to them and make choices that maximize their expected profit or expected utility. However, research has shown these assumptions are often violated. That is, people do not always behave rationally.

Behavioural economics examines the behavioural assumptions of neoclassical economics more critically. Agents are recognised to have non-rational tendencies that often contradict the assumptions and predictions of neoclassical economic theory. Incorporating work of other disciplines, primarily psychological, behavioural economics aims to move away from some of these unrealistic assumptions to more accurately describe and predict agents' economic decisions. Despite the increasing popularity of behavioural economics as a discipline, its methods and principles are applied most often in the consumer behaviour literature. The application of behavioural economics principles to agricultural producers, and the potential implications of behavioural economics in agricultural economics, have received comparatively less attention.

The three essays in dissertation take a behavioural approach to decisions of agricultural producers in risky scenarios, paying particular attention to the gambler's and hot hand fallacies (Essay 1), third generation prospect theory (Essay 2), and regret theory (Essay 3).

This work contributes to the agricultural economics literature by addressing some of the shortcomings in the neoclassical models frequently used in the discipline. By moving beyond the assumptions of these models and applying some principles of behavioural economics, I endeavour to more accurately model agents' behaviour in the face of risk. Farmers must make important and potentially consequential decisions about their operations. How farmers behave in risky situations is an important area of study, with potential implications for the natural environment, commodity supply, and policy at many levels. Studying and understanding how farmers behave in risky situations is important for how they respond to policy changes, as well as understanding the consequential decisions they make for their operations, which have the potential to more broadly impact the economy and natural environment.

ACKNOWLEDGEMENTS

I would like to thank my advisors, Dr. David Hennessy and Dr. Hongli Feng for their guidance, patience, and encouragement, and for pushing me to make this work better. Working with you both has made me a better researcher.

I would also like to thank my committee members, Dr. J. Roy Black, Dr. Robert Myers, and Dr. Eduardo Nakasone, for their invaluable feedback. Your comments and suggestions for my dissertation have been invaluable and have undoubtedly enhanced the quality of this work.

I am grateful to my friends and Peeps in the department of Food, Agricultural, and Resource Economics for their support throughout this arduous process. Thank you for your encouragement, understanding, and tolerating all the tears. I hope that I was able to provide you with a fraction of the support that you have given to me.

To my non-AFRE friends and family, your support and encouragement meant the world to me. Thank you for reminding me that there is life outside of graduate school!

Finally, the biggest thanks goes to my parents. I can say unequivocally that without your continued encouragement and support, and unwavering belief in me, I would not have made it this far.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Introduction

Neoclassical economic theory hinges on the assumption of rational behaviour. The study of agents' behaviour in the face of risk and uncertainty is often guided by this assumption: agents are assumed to make predictions based on the information available to them and make choices that maximize their expected profit or expected utility. However, research has shown these assumptions are often violated. That is, people do not always behave rationally.

Behavioural economics examines the behavioural assumptions of neoclassical economics more critically. Agents are recognised to have non-rational tendencies that often contradict the assumptions and predictions of neoclassical economic theory. Incorporating work of other disciplines, primarily psychological, behavioural economics aims to move away from some of these unrealistic assumptions to more accurately describe and predict agents' economic decisions. Despite the increasing popularity of behavioural economics as a discipline, its methods and principles are applied most often in the consumer behaviour literature. The application of behavioural economics principles to agricultural producers, and the potential implications of behavioural economics in agricultural economics, have received comparatively less attention.

The three essays in dissertation take a behavioural approach to decisions of agricultural producers in risky scenarios. Risk is endemic in agricultural production: production levels are determined in large part by the prevailing weather conditions in a particular growing season, and output prices can fluctuate significantly from year to year. Farmers must make important and potentially consequential decisions before the outcomes of these and other variables are known. How farmers behave in risky situations is an important area of study, with potential implications for the natural environment, commodity supply, and policy at many levels.

The data used in these essays comes primarily from two sources. The first is a survey conducted at small focus group meetings with farmers in the Prairie Pothole Region of North and South Dakota, designed to investigate their land use and land conversion decisions. The second is a survey of corn and soy producers in Michigan and Iowa, pertaining to their crop insurance choices. Using data obtained directly from farmers is important in understanding how they behave when faced with risk and uncertainty. Previous research has demonstrated the importance of drawing from populations of experienced professionals, rather than the general population, to examine their behaviour in a variety of settings. List (2004) and List and Haigh (2009) showed that the behaviour of experienced stock traders differed from that of undergraduate students, with experienced professionals exhibiting more rational behaviour in economic experiments. In the domain of agricultural economics, Suter and Vossler (2013) compared the behaviour of undergraduate students and dairy producers in experiments designed to examine the impact of an ambient pollution tax. Their results showed that the behaviour of the two groups differed, and that farmers' behaviour differed based on the characteristics of their actual operation. Drawing from the population of experienced agricultural producers in this work allows for more accurate representation of their economic behaviour.

My first essay uses farmers' predictions and decisions in two economic experiments to test for behaviour consistent with the gambler's and hot hand fallacies. These two fallacies describe non-rational prediction and betting behaviour. They and have been studied in a variety of hypothetical and non-hypothetical settings, but have not been examined in the context of agricultural production. Combining data from the land conversion and crop insurance experiments, I examine how previous market and weather condition outcomes influence farmers' predictions for the coming year to test for the gambler's fallacy, and how previous successful bets influence the propensity to bet in the future to test for the hot hand fallacy. My results show that farmers were less likely to predict a good year after to successive good years, consistent with the gambler's fallacy. I also find moderate support for the hot hand fallacy, with farmers more likely to bet if they had a successful bet in the previous period. However, another successful bet had no additional impact.

Essay 2 explores farmers' crop insurance decisions. Using data from a survey of farmers in Michigan and Iowa, I assess the ability of third generation prospect theory to explain farmers' stated willingness to pay (WTP) for changes in crop insurance coverage level. Previous work has shown that expected utility theory, the workhorse of agents' risky decision modelling in neoclassical economic theory, insufficiently describes farmers' crop insurance purchasing behaviour (Du et al., 2016). While expected utility allows for risk averse behaviour, it fails to account for potential aversion to losses, which is addressed in prospect theory (Kahneman and Tversky, 1979). Third generation prospect theory (Schmidt et al., 2008), allows for uncertainty in the reference choice, rather than assuming a constant reference point as in previous conceptions of prospect theory. I propose that this theoretical framework more closely follows the decision context faced by farmers when they choose among the many crop insurance products available to them.

Using farmers' stated WTP for changes in crop insurance policy coverage level, I estimate value and probability weighting function parameters. These parameter estimates suggest that third generation prospect theory is a suitable framework through which to view farmers' crop insurance choices. This framework is extended to examine how farmers might respond to proposed reductions in federal crop insurance subsidies, and suggests that farmers would elect policies with lower coverage levels if subsidies were reduced.

The third and final essay investigates the potential role of regret in farmers' land conversion decisions. High rates of grass to crop conversion have been observed in the Prairie Pothole Region of North and South Dakota in recent years. Conversion of land in this region is of concern due to the environmental services provided by area grassland, including important breeding grounds for many migratory bird species and carbon sequestration. Additionally, increased cropland may result in higher levels of fertilizer runoff into area waterways. Despite the high rates of conversion, observations suggest that conversion is lower than would be predicted if farmers made land use decisions based purely on economic considerations. I posit that anticipated regret may be a factor in landowners' observed conversion behaviour, causing them to leave land in grass despite economic incentives favouring conversion to cropland.

I explore the potential role of regret in farmers' land conversion decisions in two ways: using their stated WTP for land conversion in hypothetical scenarios, and their land conversion decisions in framed experiment. Farmers' WTP for conversion does not support the proposed theoretical framework and the hypothesis that anticipated regret influences farmers' conversion decisions. The results from the experiment, however, suggest that farmers' conversion decisions may be influenced by anticipated regret. Farmers who were made to consider the regret they might feel about their decisions converted land less frequently than those for whom regret was not made salient. Participants also stated they felt more regret about conversion decisions that non-conversion decisions (i.e. decisions to leave their land in its current use), providing support for the hypothesis that more regret is felt about decisions that change, rather than maintain, the status quo.

This work contributes to the agricultural economics literature by addressing some of the shortcomings in the neoclassical models frequently used in the discipline. By moving beyond the assumptions of these models and applying some principles of behavioural economics, I endeavour to more accurately model agents' behaviour in the face of risk. Studying and understanding how farmers behave in risky situations is important for how they respond to policy changes, as well as understanding the consequential decisions they make for their operations, which have the potential to more broadly impact the economy and natural environment.

REFERENCES

REFERENCES

Du, X., Feng, H. and Hennessy, D. A..2016. Rationality of choices in subsidized crop insurance markets. *American Journal of Agricultural Economics* 99(3): 732-756.

Kahneman, D. and Tversky, A. 1979. Prospect theory: An analysis of decision under uncertainty. *Econometrica*, 47(2): 263-293.

List, J.A. 2004. Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica* 72(2): 615-625.

List, J.A., and Haigh, M.S. 2009. Investment under uncertainty: testing the options model with professional traders. *Review of Economics and Statistics* 92(4): 974-984.

Schmidt, U., Starmer, C., and Sugden, R. 2008. Third-generation prospect theory. *Journal of Risk and Uncertainty* 36(3): 203-223.

Suter, J.F. and Vossler, C.A. 2013. Towards an understanding of the performance of ambient tax mechanisms in the field: evidence from upstate New York dairy farmers. *American Journal of Agricultural Economics* 96(1): 92-107.

CHAPTER 1. Rationality of weather predictions and insurance purchases: A test of the gambler's and hot hand fallacies based on framed experiments with farmers

Abstract

Deviations from the rational behaviour assumed in many economic models have been found in a variety of settings. Two such deviations, the gambler's and hot hand fallacies have been found in lab settings, as well as in consequential real-world decisions. Previous economic experiments have shown that the behaviour of professionals can differ from that of the general population. In this paper, we use data from two experiments conducted with a particular group of professionals who make yearly high-stakes decisions in the face of uncertain weather and market conditions: agricultural producers. In the experiments, participants were asked to make predictions about the coming year's weather and market conditions and make decisions in a familiar decision context. Results indicate evidence of the gambler's fallacy, such that participants were less likely to predict a good outcome if the previous outcome(s) were good. We also observed that participants were more likely to gamble if a previous gamble was successful, but find no impact of two successful gambles. These combined results indicate that even professionals with many years of experience can exhibit behaviours that deviate from those assumed by classical models.

1.1. Introduction

Models in neoclassical economic theory rely on rational behaviour by agents. When examining decision-making, most models in economics predict that agents make the choice that maximizes profits, or in risky situations, expected profits or expected utility. However, research has shown that people do not always make predictions according to objective probabilities and information available to them. That is, people do not always behave rationally.

How agents' perceptions of risk and risky scenarios differ from objective probabilities has been studied in several contexts (see Hansson, 2010). Two behaviours that describe deviations from rational predictions have been dubbed the hot hand and gambler's fallacies (Rabin and Vayanos, 2010). The gambler's fallacy is characterized by negative recency, such that agents believe the next event will not be the same as those preceding it (Ayton and Fischer, 2004). It is the belief that a small sequence of random outcomes should represent the underlying probabilities that generate them, rather than believing that each event will be independently determined by that underlying probability. For example, after a sequence of three heads in a fair coin toss, people who behave according to the gambler's fallacy believe that tails is more likely than heads in the next toss, since tails is 'due' (even though each toss is independent and the probability of either outcome occurring is always 0.5). This behaviour is often called belief in the law of small numbers (Tversky and Kahneman, 1971; Rabin, 2002).

The hot hand fallacy also stems from a misinterpretation of objective probabilities. It is the belief that, if a person correctly performs a task or predicts the outcome of a certain random event, that person is on a winning streak (has the hot hand) and will continue to win in the future (Ayton and Fischer, 2004). The hot hand fallacy is different from the gambler's fallacy in that it is not a belief about the outcomes *per se*, but about the person's performance (Ayton and Fischer, 2004; Croson and Sundali, 2005; Guryan and Kearney, 2008; Sundali and Croson, 2006).

Although not the opposite of the gambler's fallacy, the hot hand fallacy is typified by positive recency in that the outcome of an event such as winning a gamble is believed to be the same as the previous one.

Both concepts have been studied in a variety of contexts. Research has been conducted on subjects in a lab setting, finding evidence of both the gambler's and hot hand fallacies. Ayton and Fischer (2004) conducted experiments with undergraduate students, asking them to predict or bet on the outcome of a computerized roulette wheel. They found evidence of the gambler's fallacy among students who predicted the outcome, as they were less likely to predict a random outcome (red or blue on the computerized roulette wheel) after it had occurred. Evidence of the gambler's fallacy has also been found in people's real world decisions. Clotfelter and Cook (1993) and Terrell (1994) found that lottery numbers were less likely to be played after being drawn, inferring that individuals believed that these numbers were less likely to be drawn again, and were therefore less likely to play them. Croson and Sundali (2005) examined gambling decisions in casinos. Using video of roulette wheels to observe betting after streaks of a particular colour or number, they looked for behaviour consistent with the gambler's and hot hand fallacies. They found that people were less likely to bet on an outcome after it occurred several times (e.g., after a streak of six red outcomes, people were more likely to bet on black than after a streak of two red outcomes), consistent with the gambler's fallacy.

Chen, et al. (2016) applied the concept of the gambler's fallacy to the decisions of baseball umpires, asylum court judges, and loan applications using observational and experimental data. Controlling for pitch and game characteristics, umpires were found to be less likely to call a strike if a strike had been called on the previous pitch. Similarly, judges were less likely to grant asylum in a particular case if they had granted asylum in their previous cases, even when controlling for judge and case characteristics. In the same paper, loan officers participating in hypothetical loan approval experiments using existing loan applications were less likely to approve a loan if they had approved the previous application. This result held regardless of the compensation structure (i.e., whether loan officers were paid a flat rate or per correct approval), although the authors found that loan officers were less likely to behave according to the gambler's fallacy when they were rewarded for correct approval decisions. In all three cases (umpires, judges, and loan officers), less experienced individuals were more likely to exhibit this behaviour.

How betting behaviour is influenced by the outcomes of previous bets has also been studied in a variety of contexts. Experimental evidence of the hot hand fallacy was found among undergraduate students who placed bets on the outcome of a simulated roulette wheel, with students who had placed successful bets found more likely to bet in ensuing rounds of the experiment (Ayton and Fischer, 2004). Similarly, in an experimental setting, Rivers and Arvai (2007) found that subjects who had experienced several sequential wins were more likely to take risks in subsequent experiments than those who had experienced several losses in a row. Support for the hot hand fallacy has also been found outside the laboratory. Gilovich, Vallone, and Tversky (1985) found that basketball spectators and players were more likely to predict a successful shot when the previous shot had been successful, despite evidence to the contrary. Also examining lottery ticket purchases, Guryan and Kearney (2008) found that ticket sales increased in stores that had previously sold a winning ticket. They inferred that people viewed the stores as hot and were more willing to purchase tickets there than at non-winning stores.

List and Haigh (2009) demonstrated the importance of studying behaviours among a population of interest rather than extrapolating results from the general population (or, in their case, the population of undergraduate students). Experienced professionals will have a deeper understanding of the decision context and are more keenly aware of the consequences of different outcomes. Their responses may therefore differ from those of the general population or from undergraduate students, who are commonly used as subjects in economic experiments. To our knowledge, the gambler's fallacy and the hot hand fallacy have not been studied with a unique group of professionals who make major annual decisions in the face of uncertain weather and market conditions, such as agricultural producers. The decisions made by agricultural producers are potentially consequential: farm operations that account for the majority of agricultural production in the United States have sales of over $250,000 annually, with almost one quarter of farmland held in farms with over $1,000,000 in sales (USDA, 2017).

Risk is prevalent in agricultural production. Much of what determines productivity and farm profit are unknown to farmers when they make decisions for the upcoming growing season, such as weather outcomes and market conditions, including output and some input prices. Tools are available to farmers, especially those in developed countries like the United States, to mitigate some risk, including crop insurance, investment in certain technologies, and forward contracts, but farmers' production and income are still subject to weather and the market variations. When planning their farm activities, farmers must make predictions about conditions that will prevail and make many decisions for the coming year before the actual conditions are known. For example, they have to decide how to use their land, and what type of crop insurance, if any, to buy. As such, farmers and their unique decision contexts provide an ideal setting to examine the extent to which the gambler's and hot hand fallacies impact the predictions and decisions of experienced professionals.

In this paper, we use weather and market condition predictions in framed experiments to examine whether and how farmers' predictions deviate from rational behaviour, looking for evidence of the gambler's fallacy. Menapace, Colson, and Raffaelli (2016) showed that, among a sample of farmers, questions framed in the context of agricultural production were best able to explain farmers' actual hail insurance purchase decisions. Our experiments are framed in two agricultural production contexts with which our participants have significant experience: land use and crop insurance purchases. We also combine predictions and decisions in experimental crop insurance scenarios to look for evidence of the hot hand fallacy by examining behaviour after a successful gamble. Conducting experiments to test for these two fallacies is necessary, as observational data that would allow for such conclusions are not available. While aggregate data of farmers' yearly crop insurance purchasing behaviour exist, individual data that follows particular farmers from year to year do not. The experiments also allow us to investigate the effect of differential information dissemination, looking at how differences in communicating weather probabilities impact subjects' predictions. Our results provide limited support for the gambler's and hot hand fallacies: we observe that participants were less likely to predict an outcome after it had occurred in the past. We also see that participants were more likely to "gamble" if previous "gambles" had paid off. Our discussion poses some possible reasons for this limited support.

1.2. Conceptual framework

1.2.1. Gambler's fallacy

Formally, the gambler's fallacy can be represented by the following theoretical framework. When a sequence of successive outcomes is independently drawn, a person with casual knowledge of statistics may confound the expectation across the sequence with a conditional expectation. Consider three independent events x_t, $t \in \{1,2,3\}$, where the probability of each is $p(x_i = U) = \pi \in (0,1)$ and $p(x_i = D) = 1 - \pi$. The ex-ante probability of outcome $\{x_1 = U, x_2 = U, x_3 = U\}$ is $p(U,U,U) = \pi^3$ whereas the time $t = 2$ probability of U at time $t = 3$ is π, regardless of prior outcomes. Thus, the probability that $x_3 = U$ given that the prior two outcomes were U is $p(U \mid U,U) = \pi$, which is larger than $p(U,U,U) = \pi^3$. The casual observer may mistakenly believe that, as the ex-ante probability of three Us is smaller than $p(U \mid U,U)$ $= \pi$, independence requires that $x_3 = D$ in order to bring the average closer to the ex-ante expectation (Clotfelter and Cook, 1993; Terrell, 1994). If he holds this belief, and if the preceding two outcomes were both U, then he will be more likely to assign a higher probability to the outcome $x_3 = D$ than the objective probability of that event occurring.

Rabin (2002) makes the analogy of independent draws from an urn; someone who behaves according to the gambler's fallacy (the law of small numbers) predicts outcomes as if the sequence is being generated by drawing out of an urn without replacement. For a sequence of length N, the agent believes that πN draws will be U, and $(1-\pi)N$ draws will be D, where πN and $(1-\pi)N$ are integers. Thus, after n outcomes have been observed, the agent's beliefs about the probabilities update to $(\pi N - U_n)/(N-n)$ and $[(1-\pi)N - D_n]/(N-n)$, where U_n and D_n are the number of U and D outcomes observed in the n realizations (Rabin, 2002).

1.2.2. Hot hand fallacy

A similar framework can be used to represent the hot hand fallacy. The hot hand fallacy is a belief that, after correctly predicting a random outcome, a person has the "hot hand" and will continue to predict correctly. Rather than a belief about the random event itself, as with the gambler's fallacy, the hot hand fallacy is a belief about the person predicting the outcome(s).

Again, consider three independent events x_t, $t \in \{1,2,3\}$, where $p(x_i = U) = \pi \in (0,1)$ and $p(x_i = D) = 1 - \pi$. The ex-ante probability of outcome $\{x_1 = U, x_2 = U, x_3 = U\}$ is $p(U,U,U) = \pi^3$ whereas the time $t = 2$ probability of U at time $t = 3$ is π regardless of prior outcomes. If an agent correctly predicts or bets on two successive U outcomes, she may believe that she is on a winning streak and be more willing to bet that $x_3 = U$, regardless of the objective probability of another U occurring. That is, successive correct predictions may make people more likely to continue to bet after successes, rather than walking away.

1.3. Experimental procedures

Our data come from two experiments conducted with farmers in the United States. The experiments pertained to two common decision-making contexts faced by farmers: land conversion and crop insurance purchases. Experiment I was framed in the context of farmers' land conversion decisions. It was conducted at focus group meetings in the Prairie Pothole Region of North and South Dakota, an area which has seen significant grassland to cropland conversion in recent years (Wright and Wimberly, 2013). Experiment II pertained to crop insurance purchase decisions, and was completed by farmers in Michigan and Iowa. Farmers who grew at least 100 acres of corn or soybean in either of the two states were eligible to participate. These participants have significant experience with insurance purchase decisions, with almost 80% of farmers in our sample having purchased crop insurance in the past five years. Average sales for participants in our experiments were between \$250,000 and \$499,999, which points to the potential impact of their yearly decisions. Farmers in both experiment had an average of over 30 years of farming experience. As such, the decision contexts were familiar to farmers who participated in our experiments.

Both experiments were conducted as a part of a broader set of data gathering efforts from farmers. During the land use experiments, we gathered information regarding land use histories and estimates of important cost items. For the crop insurance experiments, we asked about recent experience with crop insurance and willingness to pay for alternative crop insurance policies. In this analysis, we only use farmers' predictions of weather and market conditions and their decisions in framed experiments. Both experiments included multiple years to represent decisions made by farmers in a familiar, real world context. We describe the two experiments in more detail below.

1.3.1. Experiment I: Land use decisions

The first experiment was conducted in March of 2016, at four locations along the James River, three in East Central South Dakota and one in East Central North Dakota. This region was chosen due to the high rate of conversion of grassland to grow row crops that has caused serious environmental concerns (Claassen et al., 2011; Miao et al., 2016; Wright and Wimberly, 2013).

The experiment asked participants to make decisions about a plot of land with clearly defined productivity levels under alternative scenarios of weather and market conditions (for ease of reference, referred to as conditions hereafter). Due to the prevalence of conversion of grassland to cropland, farmers in this area of the Dakotas were familiar with the context of land conversion. The scenarios supposed that the land was currently in grass and farmers were faced with the decision to leave it in grass or convert it to cropland for a particular conversion cost. Two states of nature, good and bad conditions[1], were possible. Each year's revenue was determined by the farmer's land use decision as well as the random conditions for that year. If

[1] The conditions shown to participants were normal or bad. For consistency with Experiment II, we refer to them here as good or bad conditions.

the good state occurred, returns were higher for crops. If the bad state occurred, grass yielded a higher return. Farmers were provided with information about returns to land under both uses and an annual per-acre conversion cost, as well as the conditional probabilities of good and bad years (the probability of a good/bad year following a good/bad year). These probabilities varied by round, but good years were always more likely to follow good years (i.e. $p\left(good_t \mid good_{t-1}\right) > 0.5$) and bad years were always more likely to follow bad years. Each year's conditions were revealed to each participant individually by members of the research team. The outcome generation process conformed with the Markov property, i.e., only the current state is relevant when forming expectations.

Upon learning about the current conditions, participants were asked to make a prediction about the coming year's conditions and to decide whether to leave their land in grass or convert it for crop production. After these decisions were made and recorded[2], conditions were revealed individually to participants by the researchers, and net returns to land was recorded. Decisions were then made for the next year. Two to four rounds of ten years each were played with participants[3]. The returns to land, the probabilities of good/bad years, and annual conversion costs were varied across different rounds. An example decision sheet used in Experiment I is presented in the Appendix A. Farmers' total compensation for attending the two-hour meeting was based on this experiment. Participants were compensated a proportion of their total revenue in one of the rounds chosen at random. Total compensation ranged from \$50 to \$80 per participant, averaging \$72.67.

[2] Predictions, decisions, and outcomes were recorded on a single sheet of paper for each round, so that farmers were aware of all previous predictions and outcomes.

[3] The number of rounds was determined by meeting time constraints. Four rounds were planned, but we were not able to complete four rounds at all meetings.

1.3.2. Experiment II: Crop insurance decisions

The second experiment was conducted with corn and soybean farmers in Michigan and Iowa in late 2016 and early 2017. The experiment was completed either in person or online. The in-person experiment was conducted at farmers meetings including the Thumb Ag Day held in Michigan in December 2016 and the Crop Advantage Series held by Iowa State University Extension in January 2017. Participants were again asked to suppose that they had a plot of land, this time planted in corn. Yearly revenue was determined by the weather and market conditions, which again could be good or bad, and whether or not they purchased crop insurance for that year. Farmers had the option of purchasing insurance, the cost of which was the same every year in that round. At the beginning of each year, farmers were asked to make predictions about the coming year's conditions and decide whether or not to purchase insurance. Predictions and insurance decisions were recorded on a decision sheet which showed the predictions, decisions, and revenue outcomes of every past year. Participants then rolled a single die, which determined whether their growing season was good or bad.[4] If any number from one to five was rolled, the conditions for that year were good. If a six was rolled, the conditions were bad. Farmers' revenue for that year was recorded on their decision sheet, and predictions and decisions were made for the next year. This continued for seven years; another seven-year round was played in which the revenue outcomes and insurance premium were changed. An example decision sheet for Experiment II is presented in the Appendix B.

Participants were paid a base rate for their participation, plus a portion of their revenue outcome for a year in a particular round of the experiment. The round and year were both chosen

[4] Weather and market outcomes in the online survey were determined by a random number generator, choosing an integer from 1 to 6. To simulate rolling a die, if the number was between one and five, the condition was good. If the number was a six, the condition was bad. Each number had an equal probability of being generated.

at random, by rolling a single die. Participants who completed the experiment in person were paid in cash; those who completed the online version were compensated in the form of an Amazon gift card. Compensation ranged from $19 to $50, with a mean of $32.81.

1.3.3. Experiment comparison

The design of the two experiments used for this analysis was similar, although not identical. In both Experiment I and Experiment II, participants were asked to make predictions about the coming year's conditions in a decision context similar to one faced by participants every year. They were then asked to make a decision (land conversion or crop insurance) which, combined with the realized conditions, determined that year's revenue. However, the experiments also had some differences. The nature of the decisions made in the experiments differed: land use decisions (Experiment I) made on participants' actual farm operations are typically made with a time horizon longer than one year, while the real-world crop insurance decisions (Experiment II) are likely made each growing season. For this reason, we use only the predictions made in Experiment II to test for the hot hand fallacy.

The other differences between the two experiments are in the experimental procedures. While predictions about the coming year's conditions were made in both experiments, Experiment I participants were asked to choose between good and bad years, while Experiment II participants were given the option of stating they were unsure (thereby declining to make a prediction). Additionally, the probabilities of good and bad years for Experiment I were determined by conditional probabilities, shown to each participant in person by members of the research team. In Experiment II, whether the year was good or bad was determined by a roll of a die (or by a random number generator meant to simulate the roll of a die). Previous studies have shown that information salience can affect how participants make choices in risky scenarios (Van Schie et al., 1995; Bordalo et al., 2012). We will explore how this difference may impact farmers' predictions and decisions.

1.4. Empirical methods

In the economics literature, testing for the gambler's and hot hand fallacies has been done with a variety of empirical methods. Some of the research cited in this paper has tested for the gambler's fallacy using analysis of variance (ANOVA) (Ayton and Fischer, 2004; Croson and Sundali, 2005). Terrell (1994) regressed lottery payout on the last time a number was drawn (as the winnings of the New Jersey lottery featured in the paper were determined in part by the number of bets placed on a particular number). Other papers have determined whether or not people exhibit gambler's fallacy behaviour by simply comparing the proportion of responses, such as Clotfelter and Cook (1993), who compare the frequency with which lottery numbers are played before and after being drawn. Sundali and Croson (2006) estimated the probability of a particular roulette wager as a function of previous outcomes using linear probability estimation. Chen et al. (2016) determined the existence of behaviour consistent with the gambler's fallacy by looking at the probability of a particular decision as a function of the previous decision (strike calls, loan approval, and granting asylum), employing both linear and nonlinear (probit/logit) probability models.

The hot hand fallacy has similarly been tested with several empirical methods. Ayton and Fischer (2004) examined the impact of successful predictions on subsequent gambles by ANOVA. Croson and Sundali (2005) examined roulette betting behaviour before and after a successful bet, regressing the number of bets placed on whether the previous bet was won. Using the same casino data, Sundali and Croson (2006) employed a linear probability model to estimate the impact of prior winning bets on the probability that individuals continue to gamble. In their study of lottery ticket sales, Guryan and Kearney (2008) examined the number of lottery tickets sold at a particular store based on whether or not a winning ticket had been sold at that store in the previous weeks, using lottery ticket sales as their dependent variable.

In this paper, we look for evidence of the gambler's and hot hand fallacies using probabilistic measures, most closely following Chen et al. (2016). Comparisons of predictions (good and bad years) after a string of good or bad years[5] are done as a first pass test of whether farmers in our sample exhibited behaviour consistent with the gambler's fallacy. To more formally test for the gambler's fallacy we use random effect probit panel estimation to determine the impact of previous outcomes on the probability that a farmer predicts good conditions for the upcoming year. We use data from both experiments to test for the gambler's fallacy.

Whether agents behave according to the hot hand fallacy is determined by using the data from Experiment II only[6]. We define making a bet in two ways: betting on a good year is defined as predicting a good year and not purchasing crop insurance, and betting on a bad year is defined as predicting a bad year and purchasing insurance. A successful bet is thus defined as making a particular bet and experiencing the predicted conditions that year. This definition most closely follows the work of Ayton and Fischer (2004), who examine the effect of successful gambles (on the outcome of a simulated roulette wheel) on the likelihood of continuing to gamble in an experimental setting. Using Experiment II data, we initially compare betting behaviour in the

[5] We did not examine strings of bad years with Experiment II data due to the nature of the data generating process. Because of the probabilities of good and bad years occurring, strings of bad years were somewhat rare, with 47 instances of two sequential bad years, and five occurrences of three bad years in a row. Accordingly, we only considered strings of good years for Experiment II.

[6] Experiment I data were not used to test for the hot hand fallacy because of the framing of the experiment. Land conversion decisions, the context of Experiment I, are likely to be made with a longer time horizon rather than on a yearly basis. Crop insurance decisions, the context of Experiment II, are typically made on a yearly basis. We believe that this context is more suitable for testing the hot hand fallacy.

current period, and how it may differ depending on the success or failure of bets in the previous period(s). We then use random effects probit models to more formally test the impact of a successful bet on the probability that participants bet again in the coming period.

1.4.1. Empirical method: Experiment I

To test for evidence of the gambler's fallacy using data in Experiment I, we examined the impact of outcomes in previous years on the prediction for the coming year similar to the approach in Chen et al. (2016). We used random effects panel probit regressions to estimate the probability that a good[7] year was predicted when the previous year(s) was (were) good. The probability of predicting a good year is modelled as in equation (1)

$$p(good_{i,t}) = \Phi\left(\beta_0 + \beta_1 I(1_{t-1}) + \boldsymbol{\beta}'\mathbf{control}\right) \tag{1}$$

where $\Phi(\cdot)$ denotes the standard normal CDF, $good_{i,t}$ indicates that a participant i predicted that year t would be good, $I(1_{t-1})$ is an indicator variable taking on the value 1 when the previous outcome was good and 0 otherwise. A vector of control variables, $\mathbf{control}$, is also included, which contains the known probability that a good year will follow a good year and a dummy variable to control for the year. We also include farm and farmer-specific variables (farm size, education level, number of years farming) to control for differences in farmers' on-farm experience.

In this experiment, a good year was always more likely to follow a good year by construction, as described previously. Farmers should have therefore always predicted a good

[7] For ease of readability, we refer to the impact and prediction of good year(s) in this section. Regressions were also run to estimate the impact of previous bad years on the probability that a bad year was predicted. Our regression equations and hypotheses for these estimation are the same, with bad years replacing good in (1) and (2), as well as H1 and H2.

year if the previous year was good. As such, rational decision making implies $\beta_1 > 0$. If instead farmers behave according to the gambler's fallacy, our hypothesis is

H1: In equation (1), $\beta_1 < 0$ (farmers are less likely to predict a good year after a good year has occurred).

We then extended the analysis to consider outcomes in the two previous years, with the probability of predicting a good year in period t modelled as

$$p(good_{i,t}) = \Phi\left(\beta_0 + \beta_1 I(1_{t-2}, 1_{t-1}) + \beta_2 I(1_{t-2}, 0_{t-1}) + \beta_3 I(0_{t-2}, 1_{t-1}) + \boldsymbol{\beta}'\mathbf{control}\right) \qquad (2)$$

where $I(1_{t-2}, 1_{t-1}) = 1$ indicates that the previous two years were good, $I(1_{t-2}, 0_{t-1}) = 1$ indicates that the previous year was bad and the preceding year was good, and $I(0_{t-2}, 1_{t-1}) = 1$ indicates that the previous year was good and the second to last year bad. Otherwise, these indicator variables are set equal to zero.

Similar to our hypothesis for equation (1), rational decision making implies that farmers should be more likely to predict a good year if a good year had just occurred, and less likely if the previous year was bad, such that $\beta_1 > 0$, $\beta_3 > 0$, and $\beta_2 < 0$. Additional instances of the same outcome should have no impact on a farmer's prediction, so that two good outcomes in a row should not impact predictions in the current period (i.e. $\beta_1 = \beta_3$). If instead farmers behave according to the gambler's fallacy, we hypothesize

H2: In equation (2), $\beta_1 < \beta_3 < 0$ (participants are less likely to predict a good year after a good year has occurred, and even less likely to predict a good year if there have been two good years in a row).

1.4.2. Empirical method: Experiment II

Using data from Experiment II we can estimate a similar probability to test whether farmers exhibit behaviour consistent with the gambler's fallacy. As there were only 47 instances of two sequential bad outcomes and three consecutive bad years occurred only five times, we limit our analysis to consider the impact and predictions of good years only. The regression equations for estimating the probability of predicting a good year in Experiment II are as in (1) and (2) above. The variables in the **control** vector include year and round dummy variables, and farm and farmer characteristics (age, education, and total acres operated).

The probabilities of good and bad years were determined by the roll of a fair six-sided die. The probability of a good year was 5/6 and that of a bad year was 1/6; these probabilities did not change with the previous outcome or with versions and rounds. The probability of a good year occurring was therefore not included in the regressions for this experiment. Because good years were always more likely than bad years, participants should have always predicted a good year in time t regardless of the previous outcome. In equation (1), this implies that $\beta_1 = 0$ is indicative of rational behaviour (a good outcome in the previous year has no impact on the prediction for the coming year). When examining the previous two outcomes as modelled in equation (2), the coefficients should be such that $\beta_1 = \beta_2 = \beta_3 = 0$. However, if farmers behave according to the gambler's fallacy we hypothesize that participants are less likely to predict a good year after one or two good years have occurred, as stated above in H1 and H2.

To test for the hot hand fallacy we treat farmers' insurance decisions as bets, as described above, and estimate the effect of previous successful bets on the probability that a farmer bets again in the coming year. We test the effect of sequential successful bets in the preceding two years as in Chen et al. (2016) The probability of purchasing insurance in period t is modelled as in (3) and (4)

$$p(bet_{i,t}) = \Phi\left(\lambda_0 + \lambda_1 I(1_{t-1}) + \boldsymbol{\lambda}'\textbf{control}\right) \quad (3)$$

$$p\left(bet_{i,t}\right) = \Phi\left(\lambda_0 + \lambda_1 I(1_{t-2}, 1_{t-1}) + \lambda_2 I(1_{t-2}, 0_{t-1}) + \lambda_3 I(0_{t-2}, 1_{t-1}) + \boldsymbol{\lambda}'\textbf{control}\right) \quad (4)$$

where $bet_{i,t}$ is an indicator variable with the value 1 if farmer i makes a bet in year t. $I(1_{t-1})$ is an indicator taking the value 1 if the previous bet was successful. $I\left(1_{t-1}, 1_{t-2}\right)$ is an indicator taking the value 1 the previous two bets were successful, $I\left(1_{t-2}, 0_{t-1}\right) = 1$ indicates that the previous bet was unsuccessful but the one before was successful, and $I\left(0_{t-2}, 1_{t-1}\right) = 1$ indicates the opposite. The indicator variables take the value of zero in all other situations. Included in the vector **control** are year and round indicator variables, as well as farm and farmer characteristics (age, education, and total acres operated).

Our main coefficients of interest are λ_1 and λ_3 . If farmers behave rationally then the success of previous bets should have no impact on their subsequent betting decisions, and we should observe $\lambda_1 = \lambda_2 = \lambda_3 = 0$. However, if farmers behave according to the hot hand fallacy we hypothesize

H3: In equation (3), $\lambda_1 > 0$ (participants are more likely to bet in the coming round if their bet in the previous bet was successful), and

H4: In equation (4) $\lambda_1 > \lambda_3 > 0$ (two sequential successful bets should have an even greater impact on the probability of placing a bet in the coming period)

As with estimations to test for the gambler's fallacy, random effects probit panel regressions were run to test for the hot hand fallacy. We first estimate the probability of betting in year t , and then consider betting on good years only.

1.5. Results

1.5.1. Experiment I results

Seventy-six farmers participated in Experiment I. Participants made 11 predictions[8] in two to four rounds, for a total of in 2,178 prediction observations. Summary statistics for experiment participants are given in Table 1.1. Participants had been farming for 37 years on average, and the majority had at least some post-secondary education. The average number of acres operated was over 2,000. Over 59% had converted some land on their farm in the preceding ten years. The total number of good predictions after one to five consecutive good or bad years in Experiment I are shown in Tables 1.3 and 1.4. If participants behave according to the gambler's fallacy, we expect fewer predictions of good years as the number of consecutive good years increases. Data from this experiment show that fewer participants predict a good/bad year as the number of consecutive good/bad years increases: after one good year, participants predicted another good year 66.1% of the time, while good years were predicted just 60.4% after five straight good years. A similar pattern was observed after bad years, with the proportion of bad years predicted falling from 37.3% to 33.6% after one to five bad years. This provides some preliminary evidence of the gambler's fallacy.

Regression results for the probabilities of predicting good years are presented in Table 1.4. We find some support for the gambler's fallacy in farmers' predictions in Experiment I. Examining the impact of the previous year's outcome, the coefficients were not statistically significant. That is, farmers were not more likely to predict a good year after a good year had occurred, contrary to our hypothesis H1. Nor were participants making rational predictions. A

[8] While participants had ten years in which to make conversion decisions, they predicted the conditions for year zero giving 11 predictions per round.

good year was always more likely to follow a good year in this experiment, so rational farmers should have been more likely to predict a good year after a good year had occurred, but we instead observed no effect of the previous outcome on the current year's prediction. When we include the outcomes in the two previous years, farmers were less likely to predict a good year if the previous two years were good (i.e. $\beta_1 < 0$), although we reject $\beta_1 < \beta_3 < 0$ as predicted by H2. Farmers were less likely to predict that the coming year would be good if the last year was bad (as was rational, since a bad year was always more likely to follow a bad year). When we consider the impact of bad years on predicting bad conditions for the coming year, we observe a similar pattern. Table 1.5 shows no impact of one bad year on the probability that a bad year is predicted. We see a small negative effect of two consecutive bad years, and now find support for H2 (that $\beta_1 < \beta_3$).

These results hold after controlling for farmer and farm-specific characteristics, as shown in Tables 1.4 and 1.5. We also examine the predictions only from participants who changed their predictions, excluding those who made the same prediction for every year of the experiment (19 participants in 25 rounds who always predicted a good year and four participants in seven rounds who predicted bad years). Making the same prediction in every year may indicate that a participant was optimistic about the coming year's conditions (if he always predicted a good year, despite the previous year's outcome) or did not fully consider the decision context. Our results hold if we exclude these participants (see Tables 1.A1 and 1.A2 in the appendix).

1.5.2. Experiment II results

A total of 141 participants for Experiment II, completed two rounds of seven years each for a total of 1,974 condition prediction and insurance purchase observations. Summary statistics for Experiment II participants are presented in Table 1.6. The average age of study participants was 53.3 years old, and they had an average of almost 30 years farming experience. The overwhelming majority (over 97%) of participants were male. Most had at least some post-secondary education. The number of acres operated ranged from 100 to over 9,000, averaging 1,079. Mean gross farm sales were over $250,000[9]. As mentioned previously, almost 80% had purchased crop insurance in the past five years.

As shown in Table 1.7, the proportion of good predictions remains roughly constant after successive numbers of good outcomes. However, when we expand to consider those farmers who said they were unsure about the coming year's condition (i.e. they chose not to predict a good or bad year), we see that the proportion of good predictions decreases from 46.8% to 42.9% as the number of consecutive good years increases from one to five (see Table 1.8). The proportion of bad predictions decreases slightly, from 8.9% to 7.1%, while unsure responses increases from 44.3% to 50.0%. This suggests that farmers are less willing to predict a good (bad) year after several good (bad) years have occurred, despite the fact that the probability of a good (bad) year occurring has not changed, providing some evidence of the gambler's fallacy.

When we examine participants' betting behaviour (combining their predictions and insurance purchases), we find that the proportion of bets made increases with the number of successive successful bets. As shown in Table 1.9, after one successful bet, the proportion of bets made is almost 85%. As the number of successful bets increases to two and three, bets are made in 90.5% and 94.8% of the following years, respectively. This provides evidence that respondents adhere to the hot hand fallacy.

From the probit regressions of Experiment II data, we first examined the probability that

[9] Gross farm sales were captured via categorial variables (see Table 1.6), so the exact average cannot be calculated.

participants made predictions for the coming year, considering the impact of previous years' outcomes. We find that, similar to the summary data, farmers were less likely to make predictions if the previous year was good (see Table 1.10). If the last year was good, participants were approximately 15% less likely to make a prediction about the coming year, a result which holds after controlling for farmer and farm-specific characteristics. When considering the previous two outcomes, we see an additional impact of two consecutive good years. The two right hand columns in Table 1.10 indicate that farmers were approximately 35% less likely to make a prediction after two good years in a row (relative to observing two consecutive bad years). We fail to reject $\beta_1 < \beta_3 < 0$, in support of H2.

Conditional on making a prediction about the coming year's conditions (i.e. excluding those who chose "unsure"), we find no impact of a good year in the previous period on the probability that a good year was predicted (see Table 1.11). However, when we expand to the two preceding outcomes, we find that farmers were approximately 13% less likely to predict a good year after two good years in a row than after two bad years in a row (the omitted category). We again fail to reject $\beta_1 < \beta_3 < 0$ in support of H2.

These general results held when including farmer and farm-specific variables. We also examine the predictions while excluding those of participants who make the same prediction in every year of the experiment, presenting the results in the appendix (47 participants fall into this category). In Experiment II, those who always predicted a good year may have been exhibiting rational behaviour (with the given probabilities of good and bad years, good years were always more likely than bad). Also, as before, always making the same prediction may indicate that participants did not fully consider their decision. Investigating the results without these individuals allows us to detect whether those who do not always make rational predictions behave fallaciously. When we exclude participants who did not change their predictions, the impact of two consecutive good years is roughly the same on the probability that participants make a prediction about the coming year's conditions, as shown in Table 1.A3 of the appendix. Table 1.A4 shows that, conditional on making a prediction, experiencing a good year in the previous period has no impact on the probability of predicting that the coming year will be good among those who changed their predictions at least once in a round. The impact of two consecutive good years has a negative impact on the probability that a good year is predicted; we again find support for our hypothesis that $\beta_1 < \beta_3 < 0$.

To test for the hot hand fallacy, we combined farmers' predictions and insurance purchasing decisions as bets, measuring the impact of successful bets on subsequent betting decisions (page 14 discusses how we define betting behaviour in this experiment). Participants who always or never purchased insurance were excluded from this analysis (69 and 9 participants, respectively). We initially combined betting on good and bad years to determine whether or not participants made a bet, conditional on making a prediction for the coming year, and then narrowed the analysis to consider bets on good years only.

We find no impact of previous successful bets on farmers' betting behaviour in the current period. Pooling all bets (bets on good and bad years), a farmer is not more likely to make a bet if his bet in the preceding period was successful, providing no support for H3. This is true when we expand and examine the potential impact of having two successful bets in a row, as shown in Table 1.12. However, when we consider bets on good years only, a different pattern emerges. We find that, after a successful bet, farmers are more likely to bet that the coming year will be good (and correspondingly less likely to bet that the coming year will be bad). Table 11.3 shows that, after a successful bet, participants are approximately 20% more likely to bet that the coming year will be good, a result that holds after adding various control variables. This is despite the fact that probabilities of either state occurring, and thus the probability of winning a bet, did not change from year to year in these experimental scenarios. This behaviour is consistent with our hypothesis of the hot hand fallacy (H3). Despite this, we see no additional impact of two consecutive successful bets (i.e., no support for H4).

1.6. Further discussion and conclusions

The two experiments used in this paper to test for the gambler's and hot hand fallacies had several important features. We drew from a population of experienced professionals to test for non-rational behaviour among a group who make consequential decisions given uncertain weather and market conditions on a yearly basis, rather than from the general population or from a population of undergraduate students. Participants were agricultural producers with significant experience in land conversion and crop insurance decisions, the contexts of our two framed experiments. The experimental protocol differed such that it allowed us to examine the effect of different ways of revealing stochastic conditions to participants. Conducting framed experiments allowed us to test for the behaviours among agricultural producers, which is not possible with existing observational datasets. While crop insurance data are available form some sources, these data do not exist in panel form that would allow researchers to observe individual producers' purchase decisions from year to year.

Our results show that agricultural producers exhibit some behaviours consistent with the gambler's and hot hand fallacies in framed experiments, though the results from our experiments are mixed. In Experiment I, we found that one previous outcome had no impact on farmers' predictions about conditions for the coming year, inconsistent with the gambler's fallacy. However, we found that participants were less likely to predict a good year after two sequential good years, and less likely to predict a bad year after two bad years. In Experiment II, farmers were less likely to make a prediction about the coming conditions after two consecutive good years. If they did make a prediction, participants were less likely to predict that the conditions in the coming year would be good if the preceding two years were good. Farmers made predictions about the coming year's conditions as if previous outcomes make the same outcome less likely, which provides evidence of the gambler's fallacy. However, our hypotheses are not fully supported by the data.

The fact that we find limited support for the gambler's fallacy in our experiments may be due to several factors. The way in which the weather and market conditions were revealed to farmers differed in the two experiments. In Experiment I, farmers were told of the probabilities of good and bad years occurring, which varied from round to round, and each year's outcome was revealed in person by a member of the research team. In Experiment II, conditions were determined by a single die, rolled by the participant (or by clicking a button in the online version). Although we see some support for the gambler's fallacy in Experiment I results, the evidence is stronger in Experiment II, in which the probabilities of good and bad years were not explicitly presented to participants. (This can be seen by comparing Tables 1.4 and 1.10.)

The difference in the way that information was presented and revealed may have affected how participants approached their predictions and decisions in each experiment. These results may highlight the importance of different methods of communicating with decision makers, and the salience of information given in person rather than information revealed implicitly. As mentioned previously, research has shown that making some information salient to study participants can affect their choices in risky scenarios (Van Schie et al., 1995; Bordalo et al., 2012). Revealing the conditional probability of weather and market conditions in person may have made these probabilities more salient. In contrast, the independence of the outcomes of a die roll may have been more implicit, and thus considered by participants differently.

Farmers' insurance purchases indicate that they also behaved according to the hot hand fallacy in our framed experiments. When we look at the probability of betting that the coming year will be good, we find an effect of previous bets. If farmers' past gambles paid off, they were more likely to bet on a good year for the coming period. However, no impact of two sequential successes was found. This difference in betting behaviour may stem from the way in which farmers view crop insurance. They may treat not purchasing insurance as gambling on good conditions for the coming year, while viewing purchasing crop insurance as an investment and protection from risk, despite the fact that they risk losing the money spent on insurance if they experience good conditions. This may have an impact on their decisions in the experimental scenarios.

Support has been found for the gambler's and hot hand fallacies in experimental and real-world settings. Those gambling at a casino or participants in an economic study unrelated to a familiar decision context (situations in which the gamblers' and hot hand fallacies were found in other studies) may approach predictions and decisions less seriously. While money may be on the line, gambling at casinos may be viewed as entertainment rather than a source of income, causing agents to make irrational predictions and behave according to the gambler's fallacy as found by Croson and Sundali (2005). When evidence of the gambler's fallacy was found in the real world decisions of baseball umpires and asylum judges in Chen et al. (2016), these decisions did not have a direct impact on agents' income (i.e. baseball umpires were not compensated based on the number of strikes they called).

While Chen et al. (2016) did find evidence of the gambler's fallacy among loan officers' approval decisions, the effects of previous decisions were not statistically significant when the loan officers were compensated according to their success rate. This suggests that when incentives are such that agent’s decisions determine compensation, they are less likely to exhibit behaviour consistent with the gambler’s fallacy. Our results add further support to the literature that professionals are less prone to behavioural fallacies in familiar decision contexts when outcomes have significant impacts. We do find some support for the gambler’s fallacy, which was based on participants’ predictions only. However, when examining farmers’ decisions and testing for the hot hand fallacy, we do not observe strong evidence in favour of non-rational decisions. These findings may reflect the importance of participants’ predictions and resulting decisions on their farm income. Although our evidence is experimental, the scenarios were framed in decision contexts that were familiar to participants and mimicked choices they faced on a yearly basis. The fact that farmers’ annual revenue directly depends on the predictions and decisions they make every year may cause them to pay closer attention to the information available to them and to use that information rationally.

The experiments used for this analysis were framed in the context of agricultural production decisions but have broader implications for professionals in risky decision contexts. For example, many commodity trading and insurance purchasing decisions are made on a yearly basis under recurring uncertainties. Exploring how agents behave in such risky settings is important for understanding how they process information and approach potentially consequential decisions.

APPENDICES

APPENDIX 1A. Tables and figures

Table 1.1. Summary statistics for Experiment I participants.

	Mean	Median	Minimum	Maximum
Years farming	37.6	39.0	10	69
Age [a]	3.5	4	1	5
Gender (% male)	97.3%	-	-	-
Education [b]	3.21	3	2	5
Expected future years farming	14.5	10	0	60
Expect that family member will take over	73.0%	-	-	-
Total acres operated	2,086	1,350	40	21,000
Gross sales [c]	3.07	3	1	5
Converted some land on their farm	59.2%	-	-	-

[a] Age coding was as follows: '1'= 19-34, '2'= 35-49, '3'= 50-59, '4'= 60-69, '5'= 70+.
[b] Education coding was as follows: '1' = primary only, '2' = high school, '3' = some college, '4' = bachelor's degree attained, '5' = advanced degree attained.
[c] Gross farm sales coding was as follows: '1' =Under $99,000, '2' =$100,000-$249,000, '3' =$250,000-$499,999, '4' =$500,000-$999,999, '5' =$1,000,000+

Table 1.2. Proportion of good weather predictions after a run of 1 to 5 good weather outcomes (Experiment I data).

Consecutive good weather outcomes	Proportion of good predictions	Observations
1	0.661	1,220
2	0.641	860
3	0.621	596
4	0.625	395
5	0.604	278

Table 1.3. Proportion of bad weather predictions after a run of 1 to 5 bad weather outcomes (Experiment I data).

Consecutive good weather outcomes	Proportion of bad predictions	Observations
1	0.372	760
2	0.341	455
3	0.356	275
4	0.366	183
5	0.336	110

Table 1.4. Probit regression results (average partial effects reported) of predicting a good year after one or two sequential good outcomes (Experiment I)

	One lag		Two lags [a]	
	1	2	3	4
Lag good outcome (β_1)	0.005	-0.010		
	(0.023)	(0.024)		
Lag good-good (β_1)			-0.060**	-0.083***
			(0.030)	(0.032)
Lag good-bad (β_2)			-0.129***	-0.133***
			(0.038)	(0.040)
Lag bad-good (β_3)			0.008	-0.020
			(0.041)	(0.042)
p(good year)	0.148	0.180*	0.127	0.164
	(0.096)	(0.099)	(0.103)	(0.106)
Controls	No	Yes	No	Yes
Observations	1,980	1,880	1,782	1,692
Log likelihood	-1112.48	-1064.78	-1016.03	-969.56

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] A one-sided Wald test of the coefficients rejects the hypothesis that $\beta_1 < \beta_3$ with p-values of 0.974 and 0.917 for regressions 3 and 4.

Table 1.5. Probit regression results (average partial effects reported) of predicting a bad year after one or two sequential bad outcomes (Experiment I)

	One lag		Two lags [a]	
	1	2	3	4
Lag bad outcome (β_1)	-0.004	-0.017		
	(0.023)	(0.024)		
Lag bad-bad (β_1)			-0.077**	-0.092***
			(0.031)	(0.032)
Lag bad-good (β_2)			-0.053	-0.062
			(0.037)	(0.038)
Lag good-bad (β_3)			0.074**	0.057
			(0.035)	(0.036)
p(bad year)	0.775***	0.742***	0.882***	0.832***
	(0.155)	(0.161)	(0.164)	(0.170)
Controls	No	Yes	No	Yes
Observations	1,980	1,880	1,782	1,692
Log likelihood	-1101.02	-1055.63	-1002.32	-876.41

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] A one-sided Wald test of the coefficients fails to reject the hypothesis that $\beta_1 < \beta_3$ with p-values of 0.000 and 0.000 for regressions 3 and 4.

Table 1.6. Summary statistics for Experiment II participants

	Mean	Median	Minimum	Maximum
Years farming	29.7	32.0	1	60
Age (years)	53.3	56	21	81
Gender (% male)	97.7%	-	-	-
Education [a]	3.8	4	1	6
Total acres operated	1,079	784	100	9,582
Gross sales [b]	3.07	3	1	5
Proportion who have purchased crop insurance	79.4%	-	-	-

[a] Education coding was as follows: '1' = less than high school, '2' = high school, '3' = some college, no degree '4' = 2-year college degree, '5' = 4-year college degree, '6' = advanced degree.
[b] Gross farm sales. 1=Under \$99,000, 2=\$100,000-\$249,000, 3=\$250,000-\$499,999, 4=\$500,000-\$999,999, 5=\$1,000,000+

Table 1.7. Proportion of good weather predictions after a run of 1 to 5 good weather outcomes (Experiment II data, unsure responses excluded)

Consecutive good weather outcomes	Proportion of good predictions	Observations
1	0.841	779
2	0.839	522
3	0.834	350
4	0.850	214
5	0.857	119

Table 1.8. Proportion of good and bad weather predictions and unsure responses after a run of 1 to 5 good weather outcomes (Experiment II data)

Consecutive bad weather outcomes	Proportion of good predictions	Proportion of bad predictions	Proportion of unsure responses	Observations
1	0.468	0.089	0.443	1401
2	0.452	0.087	0.461	970
3	0.443	0.088	0.469	659
4	0.432	0.076	0.492	421
5	0.429	0.071	0.500	238

Table 1.9. Proportion of bets made (good or bad) after a run of 1 to 3 successful bets, conditional on making a prediction (Experiment II data)

Consecutive successful bets	Proportion of bets made	Observations
1	0.846	202
2	0.905	105
3	0.948	58

Table 1.10. Probit regression results (average partial effects reported) of making a prediction after one or two sequential good outcomes (Experiment II)

	One lag		Two lags [a]	
	1	2	3	4
Lag good outcome (β_1)	-0.152***	-0.138***		
	(0.043)	(0.042)		
Lag good-good (β_1)			-0.352***	-0.336***
			(0.103)	(0.101)
Lag good-bad (β_2)			-0.210*	-0.204*
			(0.109)	(0.107)
Lag bad-good (β_3)			-0.209*	-0.211**
			(0.110)	(0.107)
Controls	No	Yes	No	Yes
Observations	1,692	1,656	1,410	1,380
Log likelihood	-781.71	-755.07	-666.09	-643.57

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] A one-sided Wald test of the coefficients fails to reject the hypothesis that $\beta_1 < \beta_3$ with p-values of 0.003 and 0.008 for regressions 3 and 4.

Table 1.11. Probit regression results (average partial effects reported) of predicting a good year after one or two sequential good outcomes (Experiment II, unsure responses excluded)

	One lag		Two lags [a]	
	1	2	3	4
Lag good outcome (β_1)	-0.016	-0.011		
	(0.019)	(0.020)		
Lag good-good (β_1)			-0.128*	-0.132*
			(0.072)	(0.074)
Lag good-bad (β_2)			-0.120	-0.134*
			(0.074)	(0.077)
Lag bad-good (β_3)			-0.104	-0.111
			(0.071)	(0.074)
Controls	No	Yes	No	Yes
Observations	966	943	797	779
Log likelihood	-326.68	-316.13	-278.72	-267.98

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] A one-sided Wald test of the coefficients fails to reject the hypothesis that $\beta_1 < \beta_3$ with p-values of 0.003 and 0.008 for regressions 3 and 4.

Table 1.12. Probability of making a bet after 1 or 2 successful bets, conditional on making a prediction. Probit average partial effects reported (Experiment II, participants who did not change insurance purchases excluded)

	One lag		Two lags [a]	
	1	2	3	4
Successful bet in period $t-1$ (λ_1)	0.038	0.040	-0.010	0.064
	(0.071)	(0.072)	(0.165)	(0.157)
Lag success-success (λ_1)			-1.622	-1.274
			(203.9)	(63.62)
Lag success-fail (λ_2)			-1.787	-1.354
			(203.9)	(63.62)
Lag fail-success (λ_3)			-1.778	-1.413
			(203.9)	(63.62)
Controls	No	Yes	No	Yes
Observations	200	193	108	104
Log likelihood	-92.68	-85.65	-41.69	-37.40

Standard errors in parentheses

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

Table 1.13. Probability of betting on a good year (predicting a good year and not purchasing insurance) after 1 or 2 successful "bets", conditional on making a prediction. Probit average partial effects reported (Experiment II, participants who did not change insurance purchases excluded)

	One lag		Two lags	
Successful bet in period $t-1$ (λ_1)	0.206***	0.206***		
	(0.070)	(0.057)		
Lag success-success (λ_1)			0.089	0.208
			(0.110)	(0.150)
Lag success-fail (λ_2)			0.088	0.238
			(0.104)	(0.177)
Lag fail-success (λ_3)			0.017	0.0414
			(0.041)	(0.092)
Control	No	Yes	No	Yes
Observations	159	154	88	85
Log likelihood	-59.96	-52.94	-24.011	-19.421

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

APPENDIX 1B. Decision sheet example, Experiment I

Net returns of land uses in normal and bad state

Land use	Weather and market conditions	
	Normal	Bad
Crop	215$/acre	100$/acre
Grass	160$/acre	120$/acre

Probability of conditions in the coming year

If the year before was **NORMAL**	
Normal	0.8
Bad	0.2

If the year before was **BAD**	
Normal	0.3
Bad	0.7

Annual conversion costs this round: 5$/acre

Year	Choice		Condition outcome (Normal/ Bad)	Returns this period ($ per acre)**	Total returns ($ per acre)	What do you think next year's conditions will be? (normal/bad)
	Grass	Crop				
0	X					
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						

****you must subtract annual conversion costs when you convert from grass to crop in every period after conversion, including the period in which the conversion is made.**

APPENDIX 1C. Decision sheet, Experiment II

WITHOUT insurance, your revenue per acre is shown below

Good conditions (a 1, 2, 3, 4, or 5 is rolled)	Bad conditions (a 6 is rolled)
$600/acre	**$300/acre**

Revenue protection insurance is available for $50/acre. The indemnity payment for this insurance policy is $300/acre.

WITH insurance, your revenue per acre is shown below

Good conditions (a 1, 2, 3, 4, or 5 is rolled)	Bad conditions (a 6 is rolled)
$550/acre ($600 - $50 insurance premium)	**$550/acre** ($300 - $50 insurance premium + $300 indemnity payment)

At the beginning of each year, you will decide whether or not to purchase insurance. You will then throw a die to determine what your revenue outcome is for that year. Play will be repeated seven times representing seven years.

Decision and outcome sheet for scenario 1

Year	**Your guess of next year's conditions**	**Purchase insurance? (Y/N)**	**Number rolled (1,2,3,4,5, or,6)**	**Outcome ($/acre)**
1	□ Good □ Bad □ Unsure			
2	□ Good □ Bad □ Unsure			
3	□ Good □ Bad □ Unsure			
4	□ Good □ Bad □ Unsure			
5	□ Good □ Bad □ Unsure			
6	□ Good □ Bad □ Unsure			
7	□ Good □ Bad □ Unsure			

APPENDIX 1D. Results of alternative models

Table 1.A1. Probit regression results (average partial effects reported) of predicting a good year after one or two sequential good outcomes (Experiment I), participants who never changed their predictions excluded

	One lag		Two lags [a]	
	1	2	3	4
Lag good outcome (β_1)	0.001	-0.016		
	(0.025)	(0.026)		
Lag good-good (β_1)			-0.074**	-0.099***
			(0.033)	(0.034)
Lag good-bad (β_2)			-0.151***	-0.157***
			(0.041)	(0.043)
Lag bad-good (β_3)			-0.014	-0.026
			(0.044)	(0.045)
p(good year)	0.169	0.183*	0.141	0.155
	(0.107)	(0.110)	(0.114)	(0.117)
Controls	No	Yes	No	Yes
Observations	1,660	1,570	1,494	1,413
Log likelihood	-1023.17	-979.29	-928.49	-885.96

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] A one-sided Wald test of the coefficients fails to reject the hypothesis that $\beta_1 < \beta_3$ with p-values of 0.065 and 0.038 for regressions 3 and 4.

Table 1.A2. Probit regression results (average partial effects reported) of predicting a bad year after one or two sequential bad outcomes (Experiment I), participants who never changed their predictions excluded

	One lag		Two lags [a]	
	1	2	3	4
Lag bad outcome (β_1)	-0.007	-0.021		
	(0.0251)	(0.0260)		
Lag bad-bad (β_1)			-0.089***	-0.109***
			(0.033)	(0.034)
Lag bad-good (β_2)			0.084**	0.066*
			(0.037)	(0.039)
Lag good-bad (β_3)			-0.059	-0.070*
			(0.039)	(0.040)
p(bad year)	0.753***	0.747***	0.880***	0.856***
	(0.168)	(0.175)	(0.177)	(0.185)
Controls	No	Yes	No	Yes
Observations	1,660	1,570	1,494	1,413
Log likelihood	-1014.61	-971.80	-917.36	-876.41

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] A one-sided Wald test of the coefficients fails to reject the hypothesis that $\beta_1 < \beta_3$ with p-values of 0.000 and 0.000 for regressions 3 and 4.

Table 1.A3. Probit regression results (average partial effects reported) of probability of making a prediction after one or two sequential good outcomes (Experiment II, unsure responses excluded), participants who never changed their predictions excluded

	One lag		Two lags [a]	
	1	2	3	4
Lag good outcome (β_1)	-0.160***	-0.146***		
	(0.043)	(0.043)		
Lag good-good (β_1)			-0.363***	-0.349***
			(0.102)	(0.100)
Lag good-bad (β_2)			-0.214*	-0.209*
			(0.110)	(0.108)
Lag bad-good (β_3)			-0.217**	-0.220**
			(0.110)	(0.108)
Controls	No	Yes	No	Yes
Observations	1,368	1,332	1,140	1,110
Log likelihood	-732.05	-706.08	-617.75	-595.98

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] A one-sided Wald test of the coefficients fails to reject the hypothesis that $\beta_1 < \beta_3$ with p-values of 0.002 and 0.004 for regressions 3 and 4.

Table 1.A4. Probit regression results (average partial effects reported) of probability of predicting a good year after one or two sequential good outcomes (Experiment II, unsure responses excluded), participants who never changed their predictions excluded

	One lag		Two lags [a]	
	1	2	3	4
Lag good outcome (β_1)	-0.038	-0.026		
	(0.041)	(0.041)		
Lag good-good (β_1)			-0.259**	-0.241**
			(0.123)	(0.120)
Lag good-bad (β_2)			-0.212*	-0.201
			(0.128)	(0.125)
Lag bad-good (β_3)			-0.239*	-0.242*
			(0.129)	(0.126)
Controls	No	Yes	No	Yes
Observations	642	619	527	509
Log likelihood	-299.64	-289.31	-252.11	-241.62

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] A one-sided Wald test of the coefficients rejects the hypothesis that $\beta_1 < \beta_3$ with p-values of 0.184 and 0.224 for regressions 3 and 4.

REFERENCES

REFERENCES

Ayton, P., and Fischer, I. 2004. The hot hand fallacy and the gambler's fallacy: two faces of subjective randomness? *Memory & Cognition* 32(8): 1369-1378.

Bordalo, P., Gennaioli, N. and Shleifer, A., 2012. Salience theory of choice under risk. *The Quarterly journal of economics*, 127(3), 1243-1285.

Chen, D.L., Moskowitz, T.J. and Shue, K. 2016. Decision making under the gambler's fallacy: evidence from asylum judges, loan officers, and baseball umpires. *Quarterly Journal of Economics* 131(3): 1181-1242.

Claassen, R.L., Carriazo, F., Cooper, J.C., Hellerstein, D and Ueda, K. 2011. Grassland to Cropland Conversion in the Northern Plains: The Role of Crop Insurance, Commodity, and Disaster Programs. ERR-120, USDA Economic Research Service.

Clotfelter, C.T. and Cook, P.J. 1993. The gambler's fallacy in lottery play. *Management Science* 39(12): 1521-1525.

Croson, R., and Sundali, J. 2005. The gambler's fallacy and the hot hand: Empirical data from casinos. *Journal of Risk & Uncertainty* 30(3): 195-209.

Gilovich, T., Vallone, R., and Tversky, A. 1985. The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, *17*(3):295-314.

Guryan, J. and Kearney, M.S. 2008. Gambling at lucky stores: empirical evidence from state lottery sales. *American Economic Review* 98(1): 458-473.

Hansson, S.O. 2010. Risk: objective or subjective, facts or values. *Journal of Risk Research* 13(2), 231-238.

List, J.A. 2004. Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica* 72(2): 615-625.

List, J.A., and Haigh, M.S. 2009. Investment under uncertainty: testing the options model with professional traders. *Review of Economics and Statistics* 92(4): 974-984.

Menapace, L., Colson, G., and Raffaelli, R. 2016. A comparison of hypothetical risk attitude elicitation instruments for explaining farmer crop insurance purchases. *European Review of Agricultural Economics.* 43(1): 113-135.

Miao, R., Hennessy, D.A., and Feng, H. 2014. Sodbusting, crop insurance, and sunk conversion costs. *Land Economics* 90(4): 601-22.

Rabin, M. 2002. Inference by believers in the law of small numbers. *Quarterly Journal of Economics* 177(3): 775-816.

Rabin, M., and Vayanos, D. 2010. The gambler's and hot hand fallacies: theory and applications. *Review of Economic Studies* 77: 730-778.

Rivers, L. and Arvai, J. 2007. Win some, lose some: the effect of chronic losses on decision making under risk. *Journal of Risk Research.* 10(8): 1085-1099.

Sundali, J., and Croson, R. 2006. Biases in casino betting: The hot hand and the gambler's fallacy. *Judgement and Decision Making* 1(1): 1-12.

Terrell, D., and Cook, P.J. 1994. A test of the gambler's fallacy: evidence from pari-mutuel games. *Journal of Risk & Uncertainty* 8(3): 301-317.

Tversky, A., and Kahneman, D. 1971. Belief in the law of small numbers. *Psychological Bulletin* 76(2): 105-110.

United States Department of Agriculture (USDA) National Agricultural Statistics Service. 2017. Farms and Land in Farms 2016 summary. Accessed from https://usda.mannlib.cornell.edu/usda/current/FarmLandIn/FarmLandIn-02-17-2017.pdf

Van Schie, E.C. and Van Der Pligt, J., 1995. Influencing risk preference in decision making: The effects of framing and salience. *Organizational Behavior and Human Decision Processes*, 63(3), 264-275.

Wright, C. K., and Wimberly, M. C. 2013. Recent land use change in the Western Corn Belt threatens grasslands and wetlands. *Proceeding of the National Academy of Science* 110(10): 4134-4139.

CHAPTER 2. Farmers' valuation of changes in crop insurance coverage level – a test of third generation prospect theory

Abstract

Recent work has shown that expected utility theory does not accurately characterize farmers' crop insurance purchases. Prospect theory has been proposed as a more suitable framework, allowing for loss as well as risk aversion. This work examines farmers' valuation of changes to crop insurance policies through the lens of third generation prospect theory. Rather than measure gains and losses from a static reference point, third generation prospect theory allows for uncertainty in both the reference and prospect choices, determining gains and losses on a state-by-state basis. Data were obtained from surveys of corn and soybean producers in Michigan and Iowa. Participants were asked to suppose they had a plot of land in corn with a hypothetical revenue distribution and a baseline revenue insurance policy, and stated how much they would be willing to pay or accept for insurance policies with higher or lower coverage levels. To assess the suitability of third generation prospect theory, value and probability weighting function parameters were estimated by nonlinear least squares. Parameters estimates indicate that third generation prospect theory better fits our data than prospect theory with a constant reference point. The analysis was extended to examine farmers' crop insurance responses to proposed cuts in federal crop insurance policies. This work is important for understanding how farmers value crop insurance policies and how they may respond to changes in crop insurance premiums.

2.1. Introduction

Crop insurance is an important tool that allows farmers to manage some of the risk inherent in agricultural production. In the United States, crop insurance is heavily subsidized by the federal government. Federal crop insurance subsidies were introduced in 1980 in an effort to encourage uptake and reduce disaster payments to farms by the federal government. The introduction of and increases in premium subsidies has achieved the government's goal of increasing insurance rates among farmers. The proportion of insured acres reached its peak in 2015, with 88% of all planted acres (over 210 million acres) falling slightly to 86% in 2017 (Zulauf et al., 2018). As discussed by O'Donoghue (2014) and Zulauf et al. (2018), a strong relationship exists between acres covered by federal crop insurance and the rate of subsidisation, with higher subsidisation rates encouraging adoption of higher coverage level policies.

Several crop insurance products are available to American farmers. Once they decide to insure their crop, they must choose how they want to insure their acres (basic, optional, or enterprise units) and between yield and revenue insurance. Yield insurance protects farmers from a decrease in yields only, and is paid out at the harvest price. Revenue insurance protects a farmer from a drop in revenue below his insured level, allowing for decreases in crop yield or in the price of that crop set by the Risk Management Agency (the agency that operates the Federal Crop Insurance Corporation, which manages the federal crop insurance program) (Shields, 2013). For both, farmers must choose the level of coverage to purchase which ranges from 50% to 85% of the expected value, based on their farm's past production history.

Insurance premium subsidies vary with the level of coverage that a farmer purchases. Premiums for catastrophic coverage (covering yield losses of 50% at 55% of the prevailing price) are fully subsidised by the federal government (although farmers must pay an administration fee for these insurance policies). The subsidy level for crop insurance premiums decreases with coverage level, such that those purchasing crop insurance with a higher coverage level have a smaller proportion of their insurance premiums subsidised by the government, but the actual subsidy is in fact larger than for lower levels of insurance (Shields, 2013; Du et al., 2016).

Crop insurance subsidies come at significant cost to the federal government, which subsidises an average of 62% of the premium costs of these policies (Shields, 2013). The federal government also reimburses private insurance companies for administrative costs, which totals over $1 billion annually (CBO, 2016). In 2009, approximately $5.4 billion was paid in insurance premium subsidies with over $2 billion distributed to farmers through these subsidies for corn alone (Shields, 2013). The total cost of the program in 2011 were estimated at over $11 billion, with $7.5 billion of that paid as premium subsidies (Glauber, 2013). Total costs under the current program are expected to be $88 billion between 2017 and 2026 (CBO, 2016).

Because of the significant cost of crop insurance subsidies, there have been calls for these subsidies to be reduced or eliminated. With a new federal administration in 2017 and a new Farm Bill expected in 2018, crop insurance subsidies and other supports to farmers may be reduced. The proposed 2019 Fiscal Year budget includes significant changes to crop insurance policy, limiting subsidized crop insurance eligibility to farmers with less than $500,000 adjusted gross income and reducing the mean subsidy rate from 62% to 48% (OMB, 2017). For 2017 rates, a reduction of this magnitude would save the government approximately $1 billion, although this number does not account for potential increases in disaster payments to compensated uninsured producers (Zulauf et al., 2018). How farmers respond to changes in subsidies and consequent changes in their insurance premiums is an important subject of study. Previous studies have investigated the relationship between premium price and insurance demand, finding that farmers' demand for insurance is elastic, such that farmers would likely respond to increases in premiums by reducing their coverage levels (O'Donoghue, 2014). If reductions in subsidies cause farmers to make changes to their coverage levels or insurance decisions, an increase in disaster payments may be observed in the event of significant decreases in yield or revenue. Farmers may also change their production plans if their insurance premiums increase.

Agents' insurance purchasing behaviour is typically modelled with an expected utility theoretical framework. In the face of risky outcomes, agents are assumed to be expected utility maximizers. For insurance choices, including those for crop insurance, expected utility theory predicts that farmers will choose the policy that maximizes their expected utility of profits from crop production. Expected utility theory generally posits that agents have a concave utility function to incorporate risk aversion. If insurance is available at an actuarially fair premium (such that the insurance premium is equal to the expected indemnity), risk averse agents should fully insure their losses under this theoretical framework (Mas Colell et al., 1995).

Despite the popularity of these models, however, recent research has shown that crop insurance purchase decisions are not always guided by the expected utility framework. Using data from crop insurance policies purchased in 2009 by American corn and soybean producers, Du, Feng, and Hennessy (2016) demonstrated that farmers' crop insurance choices are inconsistent with expected utility maximization. They showed that the coverage level elected by farmers was, on average, lower than the coverage level expected if farmers behaved as expected utility maximizers. Nor did farmers choose the policy level that maximized their subsidy. Contrary to subsidy maximisation, the coverage level chosen by farmers decreased with an increase in out-of-pocket prices of insurance policies (prices net of any government subsidies), despite the fact that the dollar value of subsidies increased with coverage level.

In many instances of expected utility failing to explain observed insurance behaviour, prospect theory has been suggested as an alternative, whether over- or under-insurance is observed (e.g. Du et al., 2016; Sydnor, 2010). Developed by Kahneman and Tversky (1979), prospect theory differs from expected utility theory in that it determines gains and losses with respect to a particular reference point; these gains and losses are treated differently by agents. In prospect theory the disutility of losing a certain amount relative to the reference point is greater in magnitude than the utility experienced from gaining the same amount relative to that reference point. Agents are therefore said to be loss averse. Rather than an expected utility function concave over its entire support, prospect theory posits a value function concave over gains and convex over losses leading to risk aversion in the gain domain, but risk seeking behaviour in the loss domain. (Kahneman and Tversky, 1979). Prospect theory also introduces nonlinearity in probabilities with a probability weighting function, which over-weighs low probability events and under-weights high probability events (Kahneman and Tversky, 1979).

Prospect theory has been applied to insurance purchases as an alternative to expected utility theory in several different contexts (Barberis, 2013). Examining a large number of home insurance contracts, Sydnor (2010) found that the high deductible chosen in many actual home insurance policies implied an unreasonably high level of risk aversion under expected utility theory. The probability weighting function, which overweighs low probability events, was able to explain the chosen deductibles not explained by risk aversion alone. In the context of home insurance, this implies an overweighting of low probably but potentially catastrophic events, leading homeowners to over insure from an expected utility standpoint (Sydnor, 2010). Barseghyan et al. (2013) also found evidence of loss aversion in home and auto insurance contract choice. Observing that people chose a deductible larger than that which would be predicted by expected utility theory, they found evidence of loss aversion resulting primarily from overweighing low probability events.

In a purely theoretical model, Schmidt (2015) demonstrates that, in a two-state world (the agent either experiences a loss or no loss) prospect theory prediction that agents should either purchase no insurance or fully insure (i.e. there is no interior solution). These results hold when the uninsured status quo and wealth with insurance are used as reference points. However, the specification of prospect theory employed by Schmidt induced loss aversion from the value function alone rather than the value function and probability weights.

In the context of agricultural production, prospect theory has been used in a limited amount to explain farmers' behaviour. Bocquého et al. (2014) conducted experiments with farmers in France to determine whether expected utility or cumulative prospect theory (Tversky and Kahneman, 1992) better explained farmers' decisions. Estimating prospect theory parameters of French farmers through multiple price list games developed by Tanaka et al. (2010) (similar to those developed by Holt and Laury, 2002). They found evidence of loss aversion and probability weighting, supporting the use of cumulative prospect theory models rather than those based on the expected utility framework as a model of farmer behaviour.

Liu (2013) examined adoption of a particular technology, Bt cotton, among Chinese farmers, looking at the factors that may influence adoption of the genetically modified crop. Despite the potential for higher profits by cultivating Bt cotton, some farmers were reluctant to adopt. Liu (2013) posits that this may be due to the higher cost and uncertain yield of the genetically modified cotton seed, causing farmers to experience a loss of revenue if adoption does not result in more revenue. She predicts that risk averse and loss averse farmers may therefore delay adoption. Experiments similar to those in Tanaka et al. (2010) were used to estimate prospect theory parameters, which were then used as independent variables in regressions to determine the probability of adoption. Farmers who exhibited loss aversion were likely to delay adoption of the new technology, while those whose behaviour was consistent with probability weighting (over-weighting rare events) were found to adopt earlier.

Prospect theory provides a natural theoretical lens for crop insurance choices that cannot be explained by expected utility theory. As pointed out by Du et al. (2016), it is likely that farmers have a reference outcome to which they compare yearly yield and revenue outcomes. They may be averse to revenue outcomes below this reference point. Observations of under-insurance, from an expected utility standpoint, may be due to risk-loving behaviour observed when faced with losses, due to convexity of the prospect theory value function in the loss domain. Babcock (2015) applied the prospect theory model to crop insurance choices, examining crop insurance purchases among US farmers in 2009. Using simulated crop yield and price data and accepted prospect theory parameter values, he found that the prospect theory model was better able to explain observed choices than expected utility theory. However, this finding was sensitive to the reference point used in the analysis. When insurance policies were treated as investment tools (i.e. when per-acre revenue and per-acre revenue plus out of pocket premium were used as reference points), prospect theory was not able to explain observed choices. Under prospect theory, the optimal coverage level choices were consistent with those observed in farmers' actual insurance purchases when insurance policies were treated as a standalone investment (i.e. when the reference point was defined as farmers' out of pocket premium) (Babcock, 2015).

While prospect theory has advantages over expected utility theory in explaining certain observed behaviours, in its original form uncertain outcomes are compared to a particular fixed reference point. The analysis in Babcock (2015) points to a weakness in this theoretical framework: results often depend on the choice of reference point. Under traditional specifications of prospect theory, gains and losses are typically measured with respect to an outcome observed with certainty. The results in Barseghyan et al., 2013 and Sydnor (2010) in support of prospect theory rely on using the household's expected outcome as a reference point from which gains and losses are determined. In stylized economic experiments used to measure prospect theory parameters, lotteries are most often valued with respect to a certain outcome. Some conceptualizations of prospect theory allow for stochastic reference points. For example, Kőszegi and Rabin (2006) develop a model that determines the expected utility of each outcome and uses this as the baseline against which gains and losses are determined. However, while this model allows for uncertainty in the reference point, it still assumes the same reference point in each possible state of the world.

When considering economic and agricultural events, it is unlikely to be the case that a risky prospect is compared to a certain outcome. It is possible that, since a baseline outcome may itself be risky, the way in which a farmer determines gains and losses from a particular reference point may also vary depending on the state that occurs. When deciding whether or not to purchase crop insurance, or deciding among coverage level options, farmers must compare two uncertain outcomes. This uncertainty cannot be adequately addressed in prospect theory models that assume a fixed reference point. To deal with uncertainty in the reference choice, Schmidt, Starmer, and Sugden (2008) have extended the prospect theory model. Their so-called third-generation prospect (PT^3) theory follows Sugden's (2003) rank dependent subjective expected utility framework and defines a value function using the outcome of a reference choice in the same state of the world against which gains and losses are measured. PT^3 has been shown to be consistent with WTA/WTP discrepancies in the face of uncertainty (Schmidt et al., 2008).

Despite being published several years ago, there have been few empirical tests or applications of third generation prospect theory. Sprenger (2015) conducted experiments, asking participants to choose between risky prospects, producing results consistent with third generation prospect theory. (However, his results were also consistent with the stochastic reference point proposed by Kőszegi and Rabin (2006)). In contrast, Birnbaum (2018) tested several theoretical properties of PT^3, finding the framework to be refuted by empirical data.

In this paper, we apply PT^3 to farmers' crop insurance choices. We assess the ability of PT^3 to explain farmers' valuation of changes to their crop insurance choices. Using data from surveys conducted with farmers in Michigan and Iowa, we use their reported willingness to pay (WTP) and willingness to accept (WTA) for increases and decreases in coverage level to estimate PT^3 value function parameters. We find support for PT^3 in our parameter estimations, with the parameters estimated suggesting risk and loss aversion among agricultural producers, as well as a moderate degree of probability weighting. Our estimated parameters are consistent with those estimated in other studies of agricultural producers (Bocquého et al., 2013). We also find that PT^3 parameter estimates are closer to values published in past work than those estimated with prospect theories that assume a constant reference point, providing further support for the PT^3 framework. This work furthers our understanding of how farmers chose among the crop insurance products available to them, and how they perceive production risk.

2.2. Conceptual framework

We begin by supposing that a farmer is faced with the choice of purchasing a revenue insurance policy for a unit of land on his farm for the coming growing season[10]. Let r represent his per-acre revenue, unknown when this decision is made, $\bar{r}$ his average revenue[11] (APH x price), and c his chosen coverage level. The policy will pay an indemnity if the farmer's revenue falls below his insured revenue; if his revenue is above this amount, he will receive no payment. The indemnity that a farmer will be paid is shown in (5). The fair premium (the expected value of the indemnity), is as shown in (6).

$$indemnity = \max\left\{c\bar{r} - r, 0\right\} \tag{5}$$

$$\text{fair premium} = \int_0^{c\bar{r}} \left(c\bar{r} - r\right) dF\left(r\right) \tag{6}$$

The premium paid by the farmer for the policy, p, is the value of the fair premium less the subsidy he receives. The subsidy amount $s\left(c\right)$ is determined by the coverage level, and so the subsidised premium paid by farmers for the insurance policy with coverage level ϕ is denoted by (7).

$$p\left(c\right) = (1 - s(c))\int_0^{c\bar{r}} \left(c\bar{r} - r\right) dF\left(r\right) \tag{7}$$

The farmer's per-acre revenue, w, is as shown in (8). It is determined by the revenue received for his crop, any indemnity payment he receives, and the premium he must pay for his insurance policy.

$$w\left(r, c, p\right) = \max\left[c\bar{r}, r\right] - p\left(c\right) \tag{8}$$

[10] This assumes farmers make coverage decision on a year-by-year basis, thinking only of the coming growing season.
[11] APH denotes actual production history, typically a ten-year average of historical yields used to determine premium rates.

2.2.1. Expected utility

In the expected utility framework, farmers should choose the coverage level that maximizes their expected utility of income[12], so that $c^* = \arg\max_c E\left[u\left(\max\left[c\bar{r}, r\right] - p(c)\right)\right]$, where $u(\bullet)$ is a concave utility function. An increase in coverage level increases the revenue guarantee and the probability that the farmer will receive an indemnity payment, increasing his utility, but will cost more than his original policy. For a farmer to choose a higher coverage level $c' > c^*$, his expected utility must be at least as high as his original utility.

The maximum amount that the farmer is willing to pay (WTP) for an increase in coverage level should be the amount above $p^* = p\left(c^*\right)$ that keeps his expected utility constant. That is,

$$E\left[u\left(w\left(r, c^*, p^*\right)\right)\right] = E\left[u\left(w\left(r, c', p^* + WTP_{c^* \to c'}\right)\right)\right] \tag{9}$$

Similarly, for $c'' < c^*$, the minimum amount that he should be willing to accept (his WTA) should be the amount that his expected utility is unchanged such that

$$E\left[u\left(w\left(r, c'', p''\right)\right)\right] = E\left[u\left(w\left(r, c^*, p^* - WTA_{c^* \to c''}\right)\right)\right] \tag{10}$$

For goods with close substitutes, any difference in agents' WTA and WTP will be caused only by the income effect. For increments in coverage level, this should be small, so that a farmer's WTP and WTA for changes in coverage level should not differ by much. Despite this theoretical result, previous research has consistently found that WTA exceeds WTP, often by a significant margin (Brown and Gregory, 1999; Horowitz and McConnell, 2002; Tunçel and Hammitt, 2014). This has been found with studies of physical objects, environmental quality,

[12] For simplicity, we assume zero costs of production and no income from non-farm sources.

and health, among others, and holds a variety of elicitation methods (e.g. economic experiments or hypothetical statements of WTP and WTA).

2.2.2. Prospect theory

Loss aversion is one of the proposed explanations for the observed willingness to pay/willingness to accept disparity, suggesting that people experience more disutility from a loss than utility from a gain of the same magnitude. This may explain why people are willing to pay less to obtain an item than they are willing to accept to give up that same item, as has been found in many economic experiments. Prospect theory accounts for loss aversion in a way that is not explained by expected utility theory, treating losses and gains from a particular reference point differently (Kahneman and Tversky, 1979).

For outcomes with discrete distribution functions, the expected utility framework is linear in probabilities such that the expected utility of an uncertain outcome is defined as $U(x)=\sum_{i=1}^{n}u(x_i)\theta_i$, where θ_i is the probability that state i will occur. The utility function for outcome i, $u(x_i)$, is an increasing function, concave over outcomes. In prospect theory, utility of the outcome is determined similarly, but with some key differences. The agent's value function, $V(x)$, is defined as

$$V(x)=\sum_{i=-m}^{n}v(x_i)\pi(\theta_i) \tag{11}$$

in which $v(x_i)$ is the value of x_i and $\pi(\theta_i)$ is the weighted probability of outcome i. Outcomes are defined with respect to some reference point, from which gains ($x_i > 0$) and losses ($x_i < 0$) are measured.

One of the main features of prospect theory is the way in which gains and losses are treated by agents. Gains and losses are determined with respect to the agent's particular reference point. Rather than a utility function that is concave over its entire domain (gains and losses), prospect theory posits a value function that is concave over gains but convex over losses. The magnitude of the value function may also be different for gains and losses to incorporate loss aversion observed in many scenarios, such that losses are felt more keenly than gains. The value function proposed by Kahneman and Tversky (1979) that accounts for this is shown in equation (12)

$$v(x) = \begin{cases} x^{\alpha_{gain}} & \text{if } x \geq 0 \\ -\lambda(-x)^{\alpha_{loss}} & \text{if } x < 0 \end{cases} \tag{12}$$

where $0 < \alpha_{gain}, \alpha_{loss} < 1$ (and often $\alpha_{gain} = \alpha_{loss}$ is assumed). The curvatures of the value function in the two domains are determined by α_{gain} and α_{loss}, while $\lambda > 1$ implies loss aversion.

Decision weights of probability, $\pi(\theta_i)$ in equation (11), is another way in which prospect theory differs from expected utility. Decision weights are commonly modelled such that low probability events are over weighted and high probability events are under weighted. Several weighting functions have been proposed, but the one most commonly employed is as in Kahneman and Tversky (1979). Their proposed weighing function is of the form $\pi(\theta) = \dfrac{\theta^{\beta}}{\left(\theta^{\beta} + (1-\theta)^{\beta}\right)^{\frac{1}{\beta}}}$, where β is the probability weighting parameter. This function causes the value function $V(x)$ to be non-linear in probabilities, and also contributes to observed loss aversion.

2.2.3. Cumulative prospect theory

Cumulative prospect theory, developed by Kahneman and Tversky (1992), retains the value function and decision weights of prospect theory developed earlier by the same authors (equation (12), above) (Kahneman and Tversky, 1979). However, cumulative prospect theory introduces a cumulative probability weighting function that determines decision weights for gains and losses differently, such that a prospect with n potential gains assigns any gain i , with outcomes ranked $x_i \le \ldots \le x_n$ the decision weight

$$\pi_i^+ = w^+\left(\theta_i + \ldots + \theta_n\right) - w^+\left(\theta_{i+1} + \ldots + \theta_n\right) = w^+\left(\theta_i\right) \tag{13}$$

such that $w^+\left(\theta_i + \ldots + \theta_n\right)$ is the probability of receiving at least outcome i and $w^+\left(\theta_{i+1} + \ldots + \theta_n\right)$ is the weighted probability of receiving an outcome strictly greater than i. A loss i of m total potential losses $x_m \le \ldots \le x_i$ is similarly assigned the probability weight

$$\pi^- = w^-\left(\theta_m + \ldots + \theta_i\right) - w^-\left(\theta_m + \ldots + \theta_{i-1}\right) = w^-\left(\theta_m\right) \tag{14}$$

These probability weighting functions weigh cumulative probabilities, such that the weighted probability is the weighted probability of gaining or losing at least that amount.

While cumulative prospect theory adds features to prospect theory, it still assumes a constant reference point, which may not be suitable for all decision-making contexts.

2.2.4. Third generation prospect theory

Prospect theory and cumulative prospect theory propose important alternatives to expected utility theory that may more accurately describe how agents choose among risky prospects. However, both compare risky outcomes to a certain reference point. This may not always be a reasonable assumption, especially when applying prospect theory to the context of agricultural production. Third generation prospect theory (PT^3), developed by Schmidt et al. (2008) builds on the previous versions of prospect theory, including a value function concave over gains and convex over losses as well as weighted probabilities that overweigh low probability events and underweight high probability events. However, PT^3 does not suppose a fixed reference point, and instead compares risky prospects to a reference choice that also depends on the state of nature.

The value function used in PT^3 follows the function proposed by Kahneman and Tversky (1979), and is of the form

$$v(z) = \begin{cases} z^{\alpha} & \text{if } z \geq 0 \\ -\lambda(-z)^{\alpha} & \text{if } z < 0 \end{cases} \tag{15}$$

with $0 < \alpha < 1$ indicating a function concave over gains and convex over losses, and $\lambda > 1$ indicating loss aversion. An agent's objective function is defined as

$$V(f, h) = \sum_i v(z_i)\pi(\theta_i) \tag{16}$$

where, as above, $\pi(\theta_i)$ is the weighted probability of state i occurring.

The key difference between prospect theory as proposed by Kahneman and Tversky (1979, 1992), and PT^3 is that z_i is the difference between the outcomes in state i of choice f and the reference choice h, against which gains and losses are measured, rather than a fixed reference point. The value function $v(z_i)$ is accordingly called the relative value function. In this framework, gains and losses for alternative f with respect to the reference choice h are compared for each potential state of the world are separately, such that the difference between the two outcomes in in state s_i is determined by

$$z_i = f(s_i) - h(s_i) \tag{17}$$

When $h(s_i)$ is a certain outcome, this function is equivalent to the previous conceptions of prospect theory.

In the context of crop insurance choices, we define a farmer's revenue in state i without insurance as his reference choice, h_i, and the revenue that he would receive in state i if he chose the policy with coverage level c as f_{ic}, his value of the insurance policy can be valued according to PT[3]. In each potential state of the world, the potential revenue outcomes without and with crop insurance are compared to determine whether the insurance policy results in a gain or a loss relative to his revenue without insurance. The differences in each possible state i, $z_i = f_{ic} - h_i$, are valued according to (15), and the value of the insurance policy with coverage level c, f_c, relative to revenue without insurance, the reference choice h, is determined by (16).

A farmer should choose the insurance policy with coverage level that maximizes his value function such that

$$c^* = \arg\max_c V\left(f_c, h\right) \tag{18}$$

yielding the maximized value function $V\left(f_{c*}, h\right)$.

We can also use this framework to determine how much a farmer would be willing to pay or accept for changes in his coverage level from a baseline insurance policy. The maximum amount that the farmer would be willing to pay to increase his coverage level is the amount that leaves his valuation unchanged at the maximum, such that $WTP_{c*\to c'}$ satisfies

$$V\left(f_{c*}, h\right) = V\left(f_{c'}, h, WTP_{c*\to c'}\right) \tag{19}$$

Similarly, the minimum amount that he would be willing to accept for a decrease in coverage level should be the amount such that

$$V\left(f_{c*},h\right)=V\left(f_{c''},h,WTP_{c*\to c''}\right) \tag{20}$$

In each state of the world, we define $z_{ic'} = f_{ic*} - h_i - WTP_{c*\to c'}$ and $z_{ic''} = f_{ic*} - h_i + WTA_{c*\to c''}$.

2.3. Data

Data were collected from surveys of corn and soybean farmers in Michigan and Iowa in late 2016 and early 2017. These two states were chosen to represent typical farms in the U.S. corn belt (Iowa) and states in which mixed farming is more prevalent (Michigan). Farmers who grew at least 100 acres of corn or soybeans in 2016 in either of the two states were eligible to participate. Surveys were administered to farmers through mail (77% of respondents), online (18%), and in person at farmer meetings (5%). The survey was tested in the summer of 2016. Researchers travelled to various farmer meetings in Michigan and invited attendees to complete the survey. Farmers were compensated at these meetings for their time. In late 2016 and early 2017, the researchers travelled to other meetings in Michigan and Iowa sponsored by Michigan State University and Iowa State University, respectively, at which farmers were invited to complete the survey.

The majority of surveys were completed by farmers online and through the mail in the winter and spring of 2017. Surveys were administered by the Centre for Survey Statistics & Methodology (CSSM) of Iowa State University. A sample of addresses for 2,000 farmers (1,000 in each state) was purchased from Farm Market iD and provided to CSSM staff. This sample included email addressed for approximately two thirds of these farmers. Farmers for whom email addresses were provided were initially sent letters to let them know they would receive an email with a link to the online survey. Emails were sent to 1,279 farmers (677 in Michigan and 601 in Iowa), of which 50 initially completed the online version of the survey. An additional sample file of 598 farmers, 299 in each state, was later obtained from Farm Market iD. CSSM staff prepared and mailed paper invitation letters to those respondents informing them that they would be receiving an email invitation to complete the online survey. From these additional addresses, 40 respondents completed the survey. For both samples, reminder emails were sent roughly a week after the initial electronic invitation. Respondents who completed the survey online were compensated between $19 and $28 depending on the outcome of an economic experiment not discussed in this work.

Surveys were mailed to 1,925 farmers, including those who had not completed the initial online survey and those for whom no email address was provided. Mailings included a postage paid return envelope and an incentive of $2. One week after the initial mailing, a reminder postcard was sent. An additional survey was sent to 1,531 farmers roughly three weeks after the initial mailing. A total of 470 completed surveys were returned to the CSSM. The surveys captured information about farmers' demographics, their farm operations, and past insurance choices and payments. Farmers were asked about their insurance purchase decisions and any insurance payments they received in the preceding five years (from 2011 to 2015). They were asked about other activities they employ, besides crop insurance, to mitigate risk (e.g., using futures markets, purchasing named-peril insurance policies, etc.). The survey also asked participants about the importance of non-financial factors in their insurance decisions.

To investigate how farmers value changes in coverage level from a baseline policy, they were shown a per-acre revenue distribution for corn. The hypothetical distribution was designed such that the actuarially fair insurance premium was typical for corn production in mid-Michigan. The discrete distribution indicated number of years in twenty they could expect to receive that particular revenue (see Figure 2.1). Farmers were asked to suppose that they had a revenue insurance policy with 75% coverage, with the fair premium and revenue guarantee for this policy shown. They were asked to report the maximum amount that they would be willing to pay to increase their coverage to 80% and 85%, and the minimum amount they would be willing to accept to decrease their coverage to 70% and 65%. For each insurance policy, farmers were given the average revenue and the revenue guarantee of the policy. Changes in coverage level, revenue guarantee, and the probability of making a claim from this baseline policy are given in Table 2.1. Farmers were asked to choose their WTP and WTA from given ranges. For this analysis, the mid point of each response was chosen as a farmer's WTP or WTA to evaluate the ability of third generation prospect theory to explain observed valuations. We use the data to estimate the PT^3 parameters and assess the ability of PT^3 to explain farmers' valuation of changes to the crop insurance coverage.

2.4. Empirical framework

We first examine farmers' stated WTP responses to motivate the use of prospect theory in their valuation of crop insurance policies. As mentioned preciously, with actuarially fair insurance premiums, expected utility-maximizing farmers should be willing to fully insure their losses. Accordingly, their WTP for changes in coverage level should be the same as the change in fair premium under the same conceptual framework. We also expect that farmers are equally sensitive to gains and losses when determining their WTP and WTA for changes in coverage level. If, however, farmers behave according to prospect theory, we should observe loss aversion in that they are more responsive to losses than to gains. As a first pass analysis, we determine the impact of expected losses and expected gains on their stated WTP. Gains and losses are defined with respect to the baseline 75% coverage policy in each state i. Expected loss is defined as the product of a loss in state i and the probability of sate i occurring, such that $E\left[loss_{nc}\right] = \sum_i loss_{nci} * \theta_i$, where θ_i is the probability of state i occurring. Expected gains are defined similarly, with $E\left[gain_{nc}\right]=\sum_i gain_{nci}*\theta_i$. We regress participant n's WTP on expected gains and losses, estimating

$$WTP_{nc}=\eta_0+\eta_1\mathrm{E}\left[loss_{nc}\right]+\eta_2\mathrm{E}\left[gain_{nc}\right]+\varepsilon_{nc} \tag{21}$$

where WTP_{nc} is farmer n's additional willingness to pay for the insurance policy with coverage level c, $c\in\{65\%,70\%,80\%,85\%\}$. If farmers do not exhibit behaviour consistent with prospect theory, we expect $\eta_1=\eta_2$ indicating that they are equally sensitive to gains and losses when determining their willingness to pay for the alternative coverage level. However, if farmers are loss averse, we expect $\eta_1>\eta_2$, such that they are more sensitive to losses than to gains and suggesting that prospect theory may more accurately describe their behaviour.

2.4.1. Empirical framework

After the above initial analysis, we use third generation prospect theory to examine farmers' valuation of changes to crop insurance policies, estimating the parameters of the value and probability weighting functions to assess the theoretical framework's ability to explain observed choices. The majority of studies that estimate prospect theory parameters ask participants to make binary choices between risky prospects, from which the model's parameters are estimated. The values estimated in the experiments conducted by Kahneman and Tversky (1979) are often used as a benchmark from which other parameter estimates are evaluated.[13] Rather than asking farmers to choose between to policies, we asked them to report how much they would be willing to pay or accept for policies with higher or lower coverage levels.

[13] This paper reported the median parameter values, and this method remains popular in the literature, although it has been met with some criticism. See Harrison and Swarthout (2016) for a discussion of this issue.

From a baseline of an uninsured state, let h_{ni} be the revenue that farmer n receives from his plot of land in random state i, without an insurance policy. When determining the value of a policy with coverage level c under PT[3], the farmer will compare the monetary outcome of the policy in each random state to the value he would receive if no insurance policy was purchased. We let f_{nci} represent the monetary value received by farmer n in state i with a policy that offers coverage level c, defined as

$$f_{nci} \equiv rev_{nci} - prem_c \tag{22}$$

where rev_{nci} is the revenue farmer n receives with coverage level c in state i, and $prem_c$ is the premium of that particular policy.

The monetary difference between the uninsured state and baseline without insurance in state i is defined as

$$z_{nci} \equiv f_{nci} - h_{ni} \tag{23}$$

Farmer n's valuation of a policy with coverage level c in state i is determined by

$$v\left(z_{nci}\right) = \begin{cases} z_{nci}{}^{\alpha} & z_{nci} \geq 0 \\ -\lambda\left|z_{nci}\right|^{\alpha} & z_{nci} < 0 \end{cases} \tag{24}$$

where α determines the curvature of the value function and λ determines the magnitude of loss aversion. His value of the policy with coverage level c, compared to the reference point of no insurance is determined by

$$V_{nc} = \sum_i v\left(z_{nci}\right)\pi\left(\theta_i\right) \tag{25}$$

where $\pi(\theta_i)$ is the weighted probability of being in state i. We use the same probability weighting function as in Schmidt et al. (2008), defined as $\pi(\theta)=\frac{\theta^{\beta}}{\left(\theta^{\beta}+(1-\theta)^{\beta}\right)^{\frac{1}{\beta}}}$ [14]. The parameter $\beta \in (0,1)$ determines the degree of probability weighting, such $\beta = 1$ indicates no probability weighting and the probabilities are taken at face value.

Whether a farmer would choose to be insured, based on the above framework, is determined by equation (25). If $V_{nc} > 0$, the farmer would experience a higher utility with insurance, and would therefore opt for coverage; if $V_{nc} \le 0$, he would choose to remain uninsured.

We can also use the value function to determine a farmer's valuation of crop insurance policies from a baseline coverage level to estimate value function parameters. The amount that a farmer is willing to pay (accept) for an increase (decrease) in coverage level from his baseline policy should be the amount that makes him indifferent between the two insurance policies. We define the monetary difference between the alternative and reference policy, similar to equation (23), as

$$z'_{nci} = f'_{nci} - h_{ni} \tag{26}$$

where h_{ni} is that amount that farmer n would receive in state i with the baseline coverage policy (in our scenarios, the baseline policy provided 75% revenue protection, so that $h_{ni} \equiv rev_{75i} - prem_{75}$). We define f'_{nci} as

[14] This is the probability weighting function proposed by Kahneman and Tversky (1979). Others have been proposed that retain the same qualitative properties of overweighing low probability events. See Prelec (1998), for an example.

$$f'_{nci} \equiv rev_{nci} - prem_c - WTP_{nc} \tag{27}$$

with $prem_c$ as the actuarially fair premium for the policy with alternative coverage level c. We let $WTP_{nc} > 0$ denote farmer n's maximum willingness to pay to increase his coverage level to c and $WTP_{nc} < 0$ his minimum willingness to accept to decrease his coverage level to c, $c \in \{65\%, 70\%, 80\%, 85\%\}$.

As in equations (24) and (25), let

$$v(z'_{nci}) = \begin{cases} {z'_{nci}}^{\alpha} & \text{for } z'_{nci} \geq 0 \\ -\lambda |z'_{nci}|^{\alpha} & \text{for } z'_{nci} < 0 \end{cases} \tag{28}$$

and

$$V'_{nc} = \sum_i v(z'_{nci}) \pi(\theta_i) \tag{29}$$

We can estimate parameters α and λ in the value function, and β in the probability weighting function by finding the values that equate (25) and (29). Borrowing from the random utility framework, we assume a random error term on farmers' valuation of the alternative crop insurance policy. Incorporating this error term, we suppose

$$V_{nc} = V'_{nc} + \varepsilon_{nc} \tag{30}$$

and assume that ε_{nc} has a standard normal distribution. It therefore follows that $V_{nc} - V'_{nc} \sim N(0,1)$.

We can estimate the PT[3] parameters with maximum likelihood estimation, with the likelihood function defined as

$$\mathrm{L} = \Pi_n \Pi_c \Phi(V_{nc} - V'_{nc}) \tag{31}$$

where $\Phi(\bullet)$ denotes the standard normal cumulative distribution function. The log likelihood function is therefore

$$\ln \mathrm{L}=\Sigma_n \Sigma_c \ln\left[\Phi\left(V_{nc}-V'_{nc}\right)\right] \tag{32}$$

with $\alpha, \beta, \lambda=\arg\max_{\alpha,\beta,\lambda}(\ln L)$. We also explore other model specifications, omitting the probability weighting function $\pi(\theta_i)$ (that is, assuming $\beta=1$) to estimate α and λ only. We then set $\lambda=1$ and estimate α and β. We also allow for different curvature parameters in the gain and loss domains, estimating α_{loss} and α_{gain} along with β and λ.

2.4.2. Alternative empirical framework

We propose an alternative empirical approach to estimate the value and probability weighting function parameters. This alternative approach is similar to Babcock (2015), who determines the individually optimal crop insurance coverage levels by estimating their prospect theory certainty equivalent (CE) value. The CE is the amount that agents would accept rather than an uncertain prospect or gamble; an agent is indifferent between this certain amount, valued according to her utility function, and the uncertain prospect. We take this approach to estimate PT[3] value and probability weighting function parameters, treating farmers' WTP as their certainty equivalent for a change in coverage level.

We use h_{n75i} to denote he revenue farmer n would receive with the 75% insurance policy in each possible state of nature i. This serves as the farmers' reference point in state i. The monetary value received in state i under the alternative policies is represented by f_{nci}, with

$c \in \{65\%, 70\%, 80\%, 85\%\}$ denoting the alternative coverage levels. The monetary difference between the baseline and alternative policies in state i is defined as

$$z_{nci} = f_{nci} - h_{n75i} \tag{33}$$

Farmer n's valuation in state i of a policy with coverage level c, compared to 75% coverage is determined by

$$v(z_{nci}) = \begin{cases} z_{nci}{}^{\alpha} & z_{nci} \geq 0 \\ -\lambda |z_{nci}|^{\alpha} & z_{nci} < 0 \end{cases} \tag{34}$$

Farmer n's value of a policy with coverage level c, compared to the 75% coverage policy, is

$$V_{nc} = \sum_i v(z_{nci}) \pi(\theta_i) \tag{35}$$

again using the same probability weighting function as above.

A farmers' maximum willingness to pay for a higher coverage level and his minimum willingness to accept for a lower coverage level is the amount that he would pay or accept with certainty for an uncertain gain or loss in revenue. We therefore treat this amount as his certainty equivalent (CE), valued according to his utility function.[15] Letting $WTP_{nc} > 0$ denote farmer n's maximum willingness to pay to increase his coverage level to $c \in \{80\%, 85\%\}$, and $WTP_{nc} < 0$ his minimum willingness to accept for $c \in \{65\%, 70\%\}$, WTP_{nc} should be such that

$$U(WTP_{nc}) = V_{nc} = \sum_i v(z_{nci}) \pi(\theta_i) \tag{36}$$

[15] We value the agent's certainty equivalent according to his utility function rather than his value function. Experiments conducted by Novemsky and Kahneman (2005) suggest that loss aversion is not exhibited when the loss is intended, such as making a payment, rather than when a loss results from a risky choice. We thus do not value his CE with the value function that incorporates loss aversion.

Supposing a constant relative risk aversion utility function, with $U\left(WTP_{nc}\right)=WTP_{nc}{}^{\alpha}$, the parameters α and λ in the value function, and the probability weighting function parameter β should be the values that satisfy (36).

We can estimate the PT[3] parameters with nonlinear least squares estimation, minimizing the sum of squared differences between the value of the change in coverage level and the CE of reported WTP (the sum of squared errors). Assuming $U\left(WTP_{nc}\right)=V_{nc}+\varepsilon_{nc}$, we estimate the parameters $\hat{\alpha},\hat{\beta}$ and $\hat{\lambda}$ that minimize

$$\sum\nolimits_{nc}\varepsilon_{nc}{}^{2}=\sum\nolimits_{nc}\left[U\left(WTP_{nc}\right)-V\left(z_{nci}\right)\right]^{2} \tag{37}$$

We also explore other model specifications, omitting the probability weighting function $\pi\left(\theta_i\right)$ (that is, assuming $\beta=1$) to estimate α and λ only. We then set $\lambda=1$ and estimate α and β. We also allow for different curvature parameters in the gain and loss domains, estimating α_{loss} and α_{gain} along with β and λ.

To compare PT[3] and prospect theory specifications with a constant reference point, we estimated value function parameters using farmers' stated WTP and the revenue guarantee of the baseline insurance policy as a constant reference point. Parameters were similarly estimated by nonlinear least squares methods, as in equation (37), but defining

$$z_{nci}\equiv f_{nci}-h_{n75}$$

where h_{n75} is the revenue guarantee of the 75% coverage policy (the same value in all possible states of nature). We estimated parameters with the same model specifications used to test PT[3] (estimating α,β, and λ, then exploring other model specifications by omitting the probability weighting function, setting $\lambda=1$, allowing α to differ in the gain and loss domains).

2.5. Results

2.5.1. Summary statistics

A total of 612 surveys were completed, with 43% of respondents operating farms in Michigan and 57% in Iowa. Summary statistics for survey respondents are presented in Table 2.2. Participants had an average age of approximately 58 years old, and had been farming for over 34 years, on average. Most respondents (over 97%) were male, and most had completed at least some post secondary education. The mean farm size was just under 960 acres, with the majority of participants growing corn and soy in the past year. Average gross farm sales were over $250,000 annually[16]. Over 80% of respondents had purchased insurance in the past five years, and almost 70% had made an insurance claim in the same time period.

In addition to MPCI, farmers used a variety of other risk management tools in their farm operations, as reported in Table 2.3. The most popular of these, employed by over 78% of respondents, was agriculture risk and price loss coverage (ARC and PLC, respectively), followed by forward and minimum price contracts (used by over 69%). Named peril insurance (e.g. hail insurance), was the third most popular of these other strategies, with approximately 60% of farmers reporting use. The others, in order of frequency, were the use of risk-mitigating technologies, such as drainage tile and other physical investments, futures and options markets, and supplemental coverage option (SCO).

2.5.2. Valuation of changes in coverage level

Average WTA and WTP for alternative coverage level policies, compared to the baseline policy with 75% revenue coverage, are plotted against changes in fair premium in Figure 2.2. The 45^{o} line indicates the change in fair premium, which should be the amount that farmers are

[16] Gross farm sales were captured with categorical variables, so an exact average value cannot be calculated.

willing to pay/accept for an increase/decrease in coverage level if they are risk-averse expected utility maximizers. As the figure shows, farmers' mean WTA for decreases in coverage level are closer to the change in fair premium than their mean WTP for increases, suggesting that the farmers value gains less than corresponding losses in coverage level. Table 2.4 shows the mean WTA and WTP responses for the different coverage level policies.

Farmers' sensitivity to decreases in coverage level is more formally demonstrated with a regression of their WTA and WTP responses on expected losses and expected gains of the alternative insurance policies. As shown in Table 2.5, the larger coefficient on expected losses indicates that expected losses have more impact on farmers' stated WTA than expected gains have on their WTP. This behaviour is consistent with prospect theory, providing motivation for exploring valuation of changes to crop insurance coverage level through this theoretical framework.

2.5.3. Prospect theory parameter estimation

The maximum likelihood estimation could not be employed as the model failed to converge after several painstaking attempts. Therefore, PT[3] parameters were estimated using the alternative empirical framework presented above, treating farmers' WTP and WTA as their certainty equivalent values for the alternative crop insurance policies and estimating parameters using nonlinear least squares. To account for the panel nature of the data (multiple WTP observations for each individual), bootstrap standard errors clustered at the individual level were calculated.

Results for all parameters estimates are presented in Table 2.6. Our statistically significant estimates of α , β , and λ (0.166, 0.444, and 1.646, respectively) denote significant risk aversion and probability weighting, and moderate loss aversion through the loss aversion parameter. Figure 2.3 shows the probability weighting function with a value of β set equal to 0.444. As shown in this figure, this value of β denotes considerable weighting of probabilities, with events with probabilities of approximately 0.25 and less given more weight than the actual probability that they would occur, and those with probabilities over 0.25 underweighted.

These values differ from the parameters estimated by Kahneman and Tversky (1979), which are often used as benchmark values in discussions of prospect theory. Their seminal paper estimated α of 0.88, β of 0.69, and λ of 2.25. These values denote moderate risk aversion, probability weighting, and loss aversion, respectively. The prospect theory parameters estimated by Liu (2013), (α = 0.48, β =0.69 and λ =3.47) and Bocquého et al. (2013) (α = 0.51, β =0.65 and λ =3.76) are similar to those in Kahneman and Tversky (1979). The studies by Liu (2013) and Bocquého et al. (2013) estimated prospect theory parameters among agricultural producers in China and France, respectively. Our estimated parameters are consistent with the qualitative conclusions of other estimates (risk aversion, probability weighting, and loss aversion) but our parameter estimates differ from those in previous work, suggesting a higher degree of risk aversion and probability weighting, and lower loss aversion from the loss aversion parameter.

We also estimated the PT[3] parameters with alternative model specifications, as outlined above. When β was set to one (no probability weighting), the estimates for α and λ are inconsistent with prospect theory. The estimated value of α of 0 suggests extreme risk aversion, such that agents would not be willing to taking on any risk. Additionally, the estimated value of λ of over 9 implies extreme loss aversion not observed in other prospect theory studies. This model specification therefore does not seem to be a good fit for our data.

The third column of Table 2.6 present the parameters estimated when λ was set to 1 (so that any loss aversion is a result solely of the probability weighting parameter β). The estimates of α and β are much closer to those estimated in previous prospect theory studies. The values of these parameter estimates suggest significant risk aversion and probability weighting, with an estimated value of α of 0.198 and an estimated β of 0.444, both significantly different from zero.

The fourth specification of the model estimated β and λ as before, but allowed for different values of α in the gain and loss domains of the value function. This estimation resulted in similar parameter values for α in the gain domain and β as in other models, but the estimates of α in the loss domain and λ were not statistically significant. The final model specification estimated α in the gain and loss domains as well as β , with λ set equal to one. The estimated values of α were 0.164 in the gain domain and 0.300 in the loss domain (both statistically significant), suggesting more risk aversion in the gain domain than risk seeking in the loss domain (a steeper curve over gains than losses). The estimated value of β is similar to that in the previous specifications.

Parameter estimates for prospect theory with the revenue guarantee of the 75% coverage policy (a constant reference point) are presented in Table 2.7. When the three value function parameters were estimated, the estimated value of α was not statistically significant from zero. The value of β , 0.312, was similar to the PT^3 parameter estimates and statistically significant. The estimated value of λ, however, denoted a higher level of loss aversion than PT^3, with a value of over 4. When we estimated different values of α in the gain and loss domains, we obtained similar results, with α not statistically different from zero in either domain. In this model specification, the value of β was no longer statistically significant, and the estimate of λ was consistent with the previous specification. When we set $\lambda = 1$ and estimated different values of α in the gain and loss domains as well as β, the parameter estimates were similar to those of PT^3, with estimated values of α_{gain} of 0.167, α_{loss} of 0.274, and β of 0.444, all statistically different from zero.

A comparison of the parameters estimated using the PT^3 and prospect theory model specifications suggest that PT^3 is more suitable than the model that compares uncertain prospects to a certain reference point. With the exception of the last model specification that estimated $\alpha_{gain}, \alpha_{loss}$ and β (i.e. when λ was set to 1), the estimates of α were not different from zero when a constant reference point was used (see Table 2.7). The estimated values of β are statistically significantly different from zero in some model specifications, but not all. The estimated values of λ are statistically significant and do denote a considerable degree of loss aversion. In contrast, the estimates for α and β are statistically significant in all model specifications and consistent across the different PT^3 models tested (see Table 2.6). The PT^3 parameter estimates are consistent with risk and loss aversion and are consistent with parameter values estimated in other studies (see Bocquého et al., 2014). We therefore suggest that PT^3 is a suitable framework through which to analyse farmers' valuation of crop insurance coverage levels.

2.6. Potential policy implications

Our parameter estimates suggest that third generation prospect theory can be used as a theoretical framework through which to examine farmers' crop insurance choices. In this section, we use the parameters estimated in the previous section to explore what this theoretical framework predicts about farmers' crop insurance purchases and the implications of proposed changes to policy premiums. Federal budget proposals include significant cuts to crop insurance subsidies, decreasing the average subsidy rate from 62% to 48%. This would result in increases in farmers' out of pocket premiums. It is not known to what extent these premium increases will change farmers' crop insurance choices.

To explore the potential ramifications of cuts to premium subsidies, we use the same hypothetical revenue distribution used in our WTP scenarios. We use parameters estimated from farmers' WTP and WTA responses (α =0.166, β 0.444, and λ =1.646), and calculate policy values according to (25). We first examine the scenario with no insurance as a baseline, determining the optimal coverage level (i.e., the one that maximizes the farmer's value function) as if he was making an initial insurance purchase under the current subsidy regime. We then use a baseline policy with 75% revenue coverage to explore whether an alternate coverage level would be valued more highly from this baseline insurance policy, again using current subsidy levels. Finally, we examine how proposed subsidy cuts might affect farmers' insurance purchasing behaviour under third generation prospect theory.

While the average crop insurance subsidy rate is 62%, policies that offer different coverage levels are subsidised at different rates. Policies that cover catastrophic losses (referred to CAT insurance policies, covering 50% of yield losses at 55% of the prevailing commodity price) are completely subsidised by the federal government. The rate of subsidisation decreases as the coverage level increases, with optional and basic unit policies offering 85% coverage subsidised at 38% (Du et al., 2016). We base our analysis on the current subsidy rates of optional and basic units, as the mean subsidy rate for these policies is 62%. (This differs from the mean subsidy rate for enterprise unit policies, which is currently 75% (Du et al., 2016)).

The valuation of crop insurance policies under PT^3 with uninsured revenue as the reference point are presented in the first column of Table 2.8. Although all the policy values are negative, indicating that remaining uninsured is the individually optimal choice under PT^3, the policy that has the highest valuation provides 75% revenue coverage. When we examine values of policies with varying coverage levels, using the 75% coverage as a baseline, we see that retaining the 75% coverage policy is still the policy with the highest value, as the value of policies with higher and lower coverage levels are all negative. This indicates that with this revenue distribution and current policy subsidy rates, farmers with a revenue insurance policy with 75% coverage should not make any changes to their coverage level under PT^3.

The proposed cuts to federal crop insurance subsidies does not specify whether the subsidies of all policies will be cut by the same proportion, only that average subsidies would be cut to 48% from 62%. To explore the changes in insurance policy values under PT^3, we reduced each subsidy level by 14%. Using these subsidy levels, we calculated the value of alternative coverage levels using a 75% policy subsidised at 55% (the current subsidy rate) as the reference point. As shown in the third column of the Table 2.8, in this scenario the 75% policy has the lowest value. The policy with the highest valuation is the 50% coverage level policy, indicating that farmers would optimize their value function by switching from 75% coverage to 50% coverage under PT^3.

Although these calculated valuations are for a stylized revenue distribution, they can offer some insight into how farmers might choose among the policies available to them. If they value policies according to PT^3, farmers would consider their revenue in many states of nature rather than using a fixed reference point. Under third generation prospect theory, remaining uninsured is personally optimal, as all insurance policies have a negative valuation. This is at odds with the American farming population, as the overwhelming majority of farmers elect at least some level of coverage. However, we observe that the policy with the highest value offers 75% coverage. This is closer to farmers' actual insurance purchase behaviour than is predicted by expected utility theory, which predicts that farmers should choose the policy with the highest coverage level (Du et al., 2016).

When we suppose that farmers are valuing alternative coverage levels from a baseline 75% policy and current average premium subsidies, we observe that all alternative coverage levels have negative values. Using the proposed premium subsidy cuts, from the current average of 62% to 48%, we observe that keeping the 75% coverage policy results in the lowest valuation. Under PT[3] and the distribution used, farmers would be better off by switching to any alternative coverage level than keeping their baseline policy. A change to any alternative coverage level would result in a higher value than remaining at 75% coverage, but policies with lower coverage levels are more highly valued than those with coverage above 75%. While this issue and this particular framework should be studied in more detail, our analysis suggests that farmers would be better off reducing their coverage level when faced with the proposed premium increases.

2.7. Further discussion and conclusions

Recent work has shown that expected utility theory to be inconsistent with farmers' crop insurance purchases (Du et al., 2016). Prospect theory is often posed as an alternative framework with which to examine agents' risky decisions. This framework has been applied in a limited extent to agricultural production and in the context of crop insurance purchases specifically. Previous work has found support for prospect theory among agricultural producers, with prospect theory found to perform better than expected utility theory in experimental settings (Bocquého et al., 2014; Liu, 2013). Prospect theory has also been found to out-perform expected utility theory in explaining farmers' observed crop insurance choices (Babcock, 2015). However, previous explorations of farmers' behaviour through the lens of prospect theory have used model specifications with a constant reference point from which gains and losses are determined. As discussed in the introduction, this may not a realistic assumption in agricultural production.

In this paper, we examined the ability of third generation prospect theory to explain farmers' reported valuation of increases and decreases in crop insurance coverage levels. We chose PT^3 to more accurately model risk in the reference choice. Rather than defining gains and losses from a constant reference point, PT^3 determines gains and losses from a risky baseline on a state-by-state basis. Using WTA and WTP data from hypothetical crop insurance parameters, we estimated parameters of PT^3 value functions, exploring various model specifications. The parameter estimates are different from those typically used in the economic literature (those estimated in Kahneman and Tversky, 1979), but they do suggest risk and loss aversion, as well as a moderate degree of probability weighting. The parameter estimates of PT^3 were more consistent with other estimates of prospect theory parameters than those estimated using a constant reference point (the revenue guarantee of the 75% coverage insurance policy), suggesting that third generation prospect theory more accurately describes farmers' behaviour.

Our findings on prospect theory suggest that farmers are both risk and loss averse. They also suggest that farmers apply non-identity decision weights rather than evaluating probabilities as given. Both of these results are consistent with traditional conceptualizations of prospect theory. However, our findings in support of PT^3 also suggest that farmers do not determine a loss from a single reference point as posited by prospect theory and cumulative prospect theory, and that considering losses on a state-by-state basis may be more suitable. While farmers may not consider eight potential states in their on-farm decision making as in our stylized crop insurance scenarios, they may consider more than one state (e.g., significant losses, outcomes that are approximately average, and above average yields) when comparing their current crop insurance contracts to alternatives available to them.

When looking at policy valuations under PT^3, we find that among the different coverage levels, the policy offering a 75% revenue guarantee is valued most highly. From the baseline 75% insurance policy, farmers' optimal policy choice remains unchanged, such that the value of every other coverage level is negative. Exploring the impact of proposed subsidy cuts, we find that the 75% coverage policy has the lowest valuation, indicating that farmers would be better off switching to any alternative coverage level, but that reducing coverage would be personally optimal.

Examining how farmers value crop insurance policies is important in understanding how they may respond to changes in crop insurance policies. Changes to federal agricultural funding have recently been proposed; these changes include significant reductions in crop insurance subsidy rates. These changes would cause potentially significant increases in the out-of-pocket premiums faced by farmers. It is important to study how farmers will respond to potential changes in their insurance premiums. Because of the current extent of crop insurance uptake (i.e. the majority of corn and soybean farmers already insure their acres with federally-subsidised crop insurance policies) it is important to consider farmers' valuation of changes to their policies from a baseline insurance policy, as with prospect theory.

Increases in crop insurance premiums are likely to impact farmers' decisions to insure their planted acres, and the coverage levels they choose. These choices may have downstream impacts on agricultural production in the United States which should be considered. Previous analyses on crop insurance subsidies have found that lower insurance premiums (through high subsidies) influence farmers' production practices and acreage decisions (Goodwin and Smith, 2013). While not all effects of crop insurance subsidies are positive (for example, farmers may convert marginal land for crop production, with negative environmental consequences (Miao et al. (2016)), how farmers will react to higher premiums, and the resulting impacts on domestic agricultural production should certainly be considered.

Consequences of crop insurance subsidy cuts may include farmers no longer electing to insure their acres or purchasing policies with lower coverage levels. Crop insurance subsidies were initially introduced in an effort to promote uptake and reduce government disaster payments. These goals were generally achieved. The role of potential decreases in insurance uptake and coverage levels should be considered in terms of their impacts on government outlays to compensate farmers in the event of catastrophic losses, especially since subsidy reductions are largely framed as decreasing federal spending on agricultural programs.

In our crop insurance scenarios in this analysis, we chose a revenue insurance policy with 75% coverage as a baseline policy, and eight possible states of nature. Further explorations into PT^3 could examine how farmers respond to different distributions and different baseline reference points, and the framing of the possible states of nature. These analyses could provide a more comprehensive picture of how farmers value crop insurance policies, and how they may react to future changes in the crop insurance products available to them.

APPENDIX

APPENDIX 2A. Tables and figures

Table 2.1. Changes in revenue guarantee and probability of payment from baseline crop insurance policy (75% coverage).

Variable			Baseline		
Coverage level	65%	70%	75%	80%	85%
Revenue guarantee	$393	$424	$454	$484	$514
Change in revenue guarantee from baseline policy (per acre)	-$61	-$30	-	+$30	+$30
Change in expected revenue from baseline policy (with no change in policy premium, per acre)	-$8.93	-$5.90	-	+$6.05	+$13.52
Probability of making a claim	0.10	0.10	0.20	0.20	0.30

Table 2.2. Summary statistics of survey respondents.

Variable	Full sample		Michigan		Iowa	
	Mean	Median	Mean	Median	Mean	Median
Number of years farming	34.2	36	35.00	37.00	33.58	35.00
Age (years)	57.95	59	57.81	59	58.05	59
Gender (% male)	97.67	-	97.96	-	97.47	-
Education [a]	3.37	3	3.23	3	3.47	3
Acres farmed	959	689.5	1163.0	800.0	815.9	600.0
Corn acres	451.6	300.0	472.1	300.0	437.8	300.0
Soy acres	364.8	250	423.4	250.0	333.3	250.0
Gross sales [b]	3.99	4	4.23	4	3.83	4
Purchased MPCI 2011-2015	80.2%	-	63.6%	-	91.9%	-
Received indemnity payment 2011-2015	69.3%	-	56.4%	-	75.5%	-

[a] Education coding was as follows: '1' = less than high school, '2' = high school, '3' = some college, no degree '4' = 2-year college degree, '5' = 4-year college degree, '6' = advanced degree.

[b] Gross sales coding was as follows: 1=Under $99,000, 2=$100,000-$249,000, 3=$250,000-$499,999, 4=$500,000-$999,999, 5=$1,000,000+

Table 2.3. Proportion of farmers reporting use of other risk management strategies, by risk management tool.

Risk management tool	Proportion Using Full sample	Michigan	Iowa
ACR/PLC	78.3%	67.8%	85.5%
Forward and minimum price contracts	69.4%	70.8%	66.9%
Named peril insurance	60.5%	43.9%	72.7%
Technologies	56.4%	57.1%	53.7%
Futures and option markets	36.6%	32.2%	39.1%
Other	7.8%	10.9%	6.0%
SCO	6.3%	6.6%	5.7%

Table 2.4. Mean hypothetical WTA and WTP for changes in coverage level from baseline 75% coverage.

	-10% (-$8.89)	-5% (-$5.90)	+5% (+$6.05)	+10% (+$13.53)
Mean response	-9.31	-8.39	4.69	7.04

Table 2.5. Impacts expected loss and expected gain on WTA and WTP (linear RE and FE regression)

	WTA/WTP RE	FE
E[loss]	1.165 ***	1.166***
E[gain]	0.558***	0.556 ***
Constant	0.108	0.117

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 2.6. PT[3] parameter estimates for various model specifications.

	1	2	3	4	5
α				0.168***	0.164***
(gain domain)	0.166***	0.500	0.198***	0.003	0.003
α				0.056***	0.300***
(loss domain)	0.003	-[a]	0.004	0.008	0.002
β	0.444***	1 (by construction)	0.440***	0.443***	0.444***
	0.000		0.000	0.000	0.000
λ	1.646***	1.000	1 (by construction)	2.470***	1 (by construction)
	0.006	-[a]		0.075	

Bootstrap standard errors, clustered at the participant level, in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] Unable to calculate estimate.

Table 2.7. Prospect theory parameter estimates, revenue guarantee of the 75% coverage policy used as the reference point

	1	2	3	4	5
α				0.000	0.167***
(gain domain)	0.000	-[a]	0.000	(0.020)	(0.010)
α				0.000	0.274***
(loss domain)	(0.019)	- [a]	- [a]	(0.081)	(0.006)
β	0.312***	1 (by	0.000	0.312	0.444***
	(0.014)	construction)	- [a]	(1.359)	(0.011)
λ	4.130***	- [a]	1 (by	4.130***	1 (by
	(0.067)	- [a]	construction)	(0.081	construction)

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] Unable to calculate estimate.

Table 2.8. Crop insurance policy valuations under PT[3], with various reference points, using estimated parameter values ($\alpha = 0.166, \lambda = 1.646, \beta = 0.444$).

Coverage level	Reference point		
	No insurance	75% coverage	
		Current average subsidy level	Proposed average subsidy level
85%	-6.09	-0.97	-1.67
80%	-5.91	-1.39	-1.63
75%	-5.40	0	-3.05
70%	-8.40	-0.16	-0.32
65%	-8.00	-0.20	-0.28
60%	-7.32	-0.20	-0.23
55%	-9.64	-0.21	-0.22
50%	-8.38	-0.21	-0.21

Figure 2.1. Hypothetical revenue distribution shown to farmers.

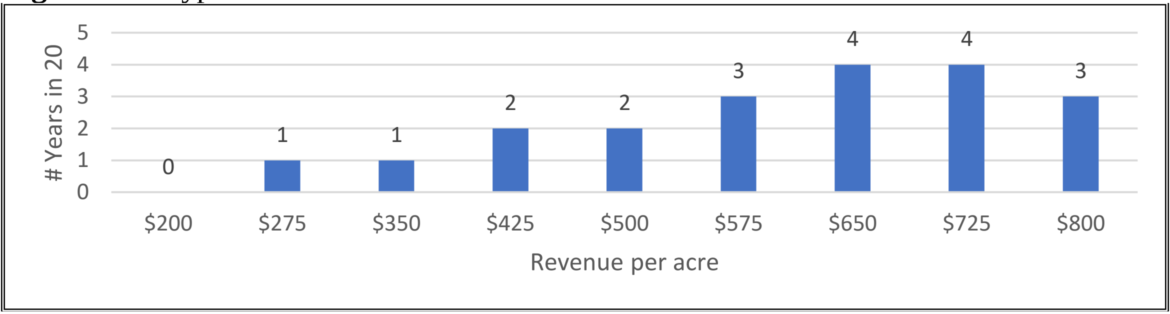

Figure 2.2. Plot of mean responses (WTA and WTP) and change in fair premium.

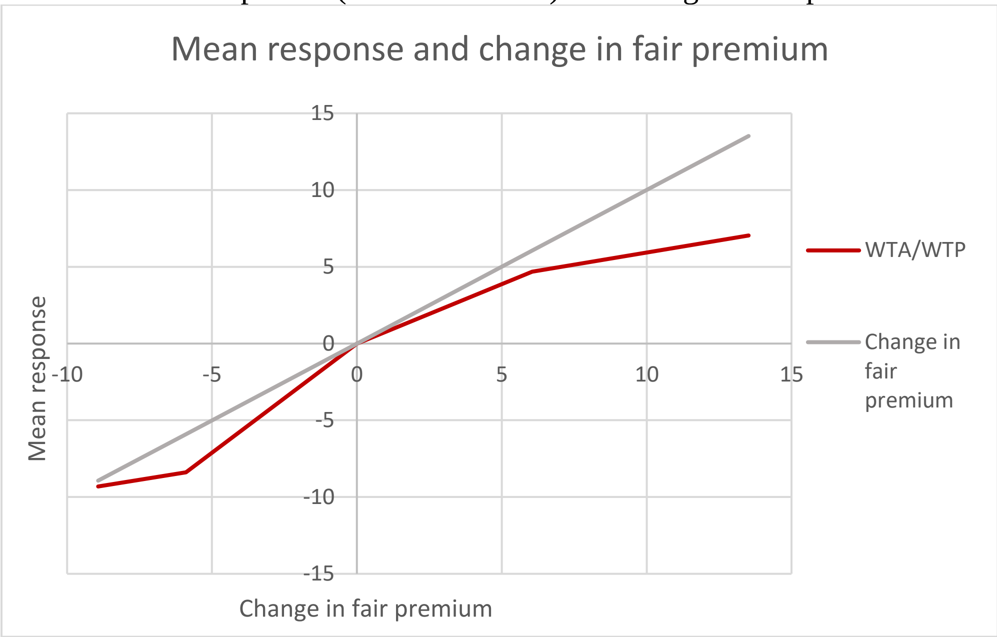

Figure 2.3. Probability weighting function with $\beta = 0.444$

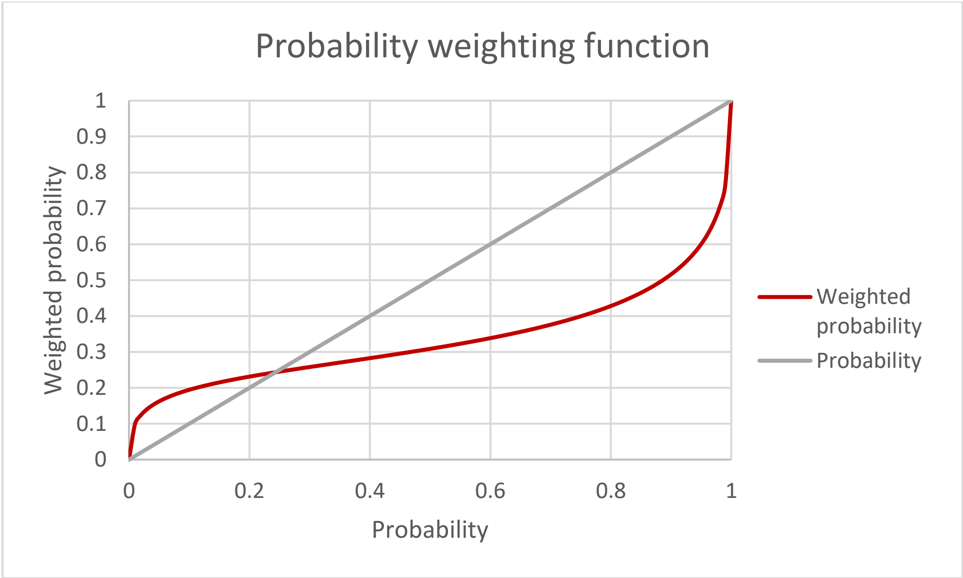

REFERENCES

REFERENCES

Babcock, B.A. 2015. Using cumulative prospect theory to explain anomalous crop insurance coverage choice. *American Journal of Agricultural Economics* 97(5): 1371-1384.

Barberis, N.C., 2013. Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives* 27(1): 173-96.

Barseghyan, L., Molinari, F., O'Donoghue, T. and Teitelbaum, J.C., 2013. The nature of risk preferences: Evidence from insurance choices. *American Economic Review* 103(6): 2499-2529.

Birnbaum, M.H. 2018. Empirical evaluation of third-generation prospect theory. *Theory and Decision* 84(1): 11-27.

Bocquého, G., Jacquet, F. and Reynaud, A. 2014. Expected utility or prospect theory maximisers? Assessing farmers' risk behaviour from field experiment data. *European Review of Agricultural Economics* 41(1): 135-172.

Brown, T.C., and Gregory, R. 1999. Why the WTP-WTA disparity matters. *Ecological Economics* 28(3): 323-332.

Congressional Budget Office (CBO). 2016. Options for Reducing the Deficit: 2017 to 2026 Reduce Subsidies in the Crop Insurance Program. Retrieved from https://www.cbo.gov/budget-options/2016/52167.

Du, X., Feng, H. and Hennessy, D. A..2016. Rationality of choices in subsidized crop insurance markets. *American Journal of Agricultural Economics* 99(3): 732-756.

Glauber, J. W. 2013. The growth of the federal crop insurance program, 1990-2011. *American Journal of Agricultural Economics* 95 (2): 482–88.

Goodwin, B. K., and Smith, V.H. 2013. What harm is done by subsidizing crop insurance? *American Journal of Agricultural Economics* 95 (2): 489–97.

Harrison, G.W. and Rutström, E.E., 2009. Expected utility theory and prospect theory: One wedding and a decent funeral. *Experimental Economics* 12(2): 133-158.

Harrison, G.W. and Swarthout, J.T., 2016. Cumulative prospect theory in the laboratory: A reconsideration. *Center for the Economic Analysis of Risk, Working Paper, 2*.

Holt, C. A., and Laury, S.K. 2002. Risk aversion and incentive effects. *American Economic Review* 92(5): 1644-1655.

Horowitz, J.K., and McConnell, K.E. 2003. A review of WTA/WTP studies. *Journal of Environmental Economics and Management* 44(3): 426-447.

Kahneman, D., and A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47(2): 263-292.

Kőszegi, B. and Rabin, M., 2006. A model of reference-dependent preferences. *The Quarterly Journal of Economics* 121(4): 1133-1165.

Liu, E.M. 2013. Time to change what to sow: Risk preferences and technology adoption decisions of cotton farmers in China. *Review of Economics and Statistics* 95(4):1386-1403.

Mas-Colell, A., Whinston, M.D., and Green, J.R. 1995. *Microeconomic theory* (Vol. 1). New York: Oxford University Press

Novemsky, N. and Kahneman, D., 2005. The boundaries of loss aversion. *Journal of Marketing Research* 42(2): 119-128.

O'Donoghue, E.J. 2014. The Effects of Premium Subsidies on Demand for Crop Insurance. ERR-169, U.S. Department of Agriculture, Economic Research Service.

Office of Management and Budget, 2017. A New Foundation for American Greatness, Budget of the US Government Fiscal Year 2018. Accessed from https://www.whitehouse.gov/omb/budget

Prelec, D., 1998. The probability weighting function. *Econometrica* 66(3):497-527.

Schmidt, U., Starmer, C., and Sugden, R. 2008. Third-generation prospect theory. *Journal of Risk and Uncertainty* 36(3): 203-223.

Schmidt, U. 2016. Insurance demand under prospect theory: a graphical analysis. *Journal of Risk and Uncertainty* 83(1): 77-89.

Shields, D. A. 2013. Federal crop insurance: Background. R40532. Congressional Research Service.

Sugden, R. 2003. Reference dependent subjective expected utility. *Journal of Economic Theory* 111(2): 172-191.

Sprenger, C. 2015. An endowment effect for risk: Experimental tests of stochastic reference points. *Journal of Political Economy* 123(6): 1456-1499.

Sydnor, J. 2010. (Over)insuring Modest Risks. *American Economic Journal: Applied Economics* 2(4): 177-199.

Tanaka, T., Camerer, C.F. and Nguyen, Q., 2010. Risk and time preferences: linking experimental and household survey data from Vietnam. *American Economic Review* 100(1): 557-71.

Tunçel, T. and Hammitt, J.K. 2014. A new meta-analysis on the WTP/WTA disparity. *Journal of Environmental Economics and Management* 68(1): 175-187.

Tversky, A., and Kahneman, D. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4): 297-323.

Zulauf, C., Schnitkey, G., Coppess, J., and Paulson, N. 2018. Premium subsidy and insured U.S. acres. *farmdoc daily* (8):46, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, March 15, 2018. Retrieved from http://farmdocdaily.illinois.edu/2018/03/premium-subsidy-and-insured-us-acres.html

CHAPTER 3. The role of regret in farmers' land conversion decisions

Abstract

Land conversion in the Prairie Pothole Region of North and South Dakota is of increasing concern. Conversion of grassland to cropland in the area has many environmental consequences, including the loss of important migratory bird breeding grounds, increased input demand, and release of sequestered carbon into the atmosphere. While conversion has negative ecological consequences, farmers face economic incentives for conversion, with cropland yielding higher returns than grassland. However. recent research suggests that famers do not always make economically optimal conversion decisions, such that they do not convert land when there are economic incentives to do so. In this paper we propose regret as a reason for under-conversion and explore regret theory in the context of land conversion in the Prairie Pothole Region. We examined the role of regret in farmers' land conversion decisions through their reported willingness to pay (WTP) for conversion in hypothetical land use scenarios. We also conducted an experiment in which farmers decided whether or not to convert their land from grass to cropland. Regret was made salient for approximately half of study participants. We observe no support for our proposed theoretical framework in farmers' WTP for conversion. However, our experimental results suggest a role of regret in farmers' land conversion decisions. Those for whom regret was made salient made fewer conversion decisions than those for whom it was not, indicating that a consideration of their future feelings of their conversion decisions may discourage land conversion. We also find that farmers made fewer decisions to convert land when cropping was the regret-maximizing choice. These experimental results suggest that regret may play a role in farmers' land conversion decisions, and that encouraging farmers to consider how they might feel about their decisions in the future may lead to more land being left in grass.

3.1. Introduction

Transitivity is among the most fundamental axioms of economic theory. Preference for A over B and for B over C should surely imply preference for A over C. And yet, when posed differently, for many this assumption does not fit well with our intuitive sense of how humans frame and make choices. A large literature has emerged documenting and seeking to explain preference reversals (Tversky and Thaler, 1990), inconsistent time preferences (Malhotra et al., 2002) and related phenomena. Among these, regret theory (e.g., Loomes and Sugden, 1982) has appeal. The theory posits a preference structure that may violate transitivity because payoffs enter both as opportunities for material consumption and as a reference point from which to measure the cost of lost opportunities (Loomes et al., 1991).

Regret may be felt if a chosen course of action is proven to be suboptimal ex post. The potential for regret exists if an agent is faced with multiple actions that can be taken and multiple states of nature are possible. If she makes the correct choice, such that her decision yielded the highest possible outcome in the realized state of nature, she will experience no regret. If, however, another choice would have yielded a higher outcome, she may regret her chosen action[17] (Loomes and Sugden, 1982). The regret a person feels will depend on the difference between the outcome from her chosen action and the maximum possible outcome in the prevailing state of nature.

In real-world decision making, regret theory has been applied to financial decision making (Michenaud and Solnik, 2008; Li et al., 2012), stock market participation and choice (Bailey and Kinerson, 2005; Fogel and Berry, 2006), consumer choice (Inman and Zeelenberg, 2002), and route choice in the transportation literature (Ben-Elia, Ishaq, and Shiftan, 2013). In

[17] Regret is felt in relation to an agent's action. In this way, regret differs from disappointment, which is felt in relation to the state of nature that has occurred, rather than to the action taken by the agent (Zeelenberg and Pieters, 2007).

their regret-theoretic expected utility framework discussed above, Braun and Muermann (2004) applied regret theory to insurance purchases. Filiz-Ozbay and Ozbay (2007) explored the role of regret in auction bids, finding experimental evidence that the anticipated regret of losing a bid can lead to bidding more than one's valuation. Ratan and Wen (2016), however, found no impact of regret on bidding behaviour.

In this paper, we explore the role of regret in land use decisions among experienced agricultural producers. Our study participants reside in Eastern South Dakota and Southeastern North Dakota, an area encompassed by the Prairie Pothole Region. Farmers in the Prairie Pothole Region are often faced with the decision of converting grassland (land in grass that has never been cropped or land previously in production but left in grass for many years) to land suitable for crop production. Technological innovations, changes in climate and weather patterns, and high commodity prices in the recent past have made row crop production more feasible in the region, promoting the conversion of grassland to land for cropping (Johnston, 2014; Wright and Wimberly, 2013). Crop production may have the potential for higher returns, but considerable conversion costs must be incurred by farmers to convert land. Additionally, fluctuating commodity prices are likely to add uncertainty to returns from crop production. Land conversion in the Prairie Pothole Region is a concern as area grassland provides important feeding and nesting ground for many migratory bird species (Claassen et al., 2011). Conversion of less productive marginal land will require more intensive input use, increasing chemical runoff into the Mississippi River watershed.

In the non-economic literature, such as the geographic and environmental sciences, land conversion is often modelled with the parcel of land, rather than the decision maker, as the unit of inquiry. Spatial models have been developed by researchers in these disciplines to examine land conversion. These models use the physical characteristics of land parcels, as well as the characteristics and uses of neighbouring parcels, to predict land use changes or to identify land most likely to be converted (Irwin and Geoghegan, 2001). Though these models may include socioeconomic factors (such as income, population density, etc.), they are not driven by economic models and thus often neglect the motivations of those making land use change decisions.

Land conversion has also been an area of study economics literature, which makes the decision maker, rather than the parcel of land, the unit of analysis (Irwin and Geoghegan, 2001). Economic theory has been applied to develop testable models of conversion, with expected profit and expected utility maximization most often guiding researchers to consider landowners' motivation(s) for their conversion decisions (Chomitz and Gray, 1996; Cohn et al., 2016; Pfaff, 1999). As an alternative to expected profit and utility maximization, real options frameworks have also been applied to land conversion decisions. This method approaches land use change from an economic perspective (i.e. profit or utility maximization) but allows for strategic delay in conversion decisions. Song et al. (2011) used real options to model the conversion of agricultural land in a corn/soy rotation to the production of biofuels, showing that the real options framework predicts different conversion patterns than optimal land use based on the net present value. Miao et al. (2015) employ a real options framework to model the conversion of grassland to land used for agricultural production, showing that risk mitigation policies may promote conversion of land out of grassland. Shah and Ando (2016) applied a real options framework to study conversion of land to agricultural and urban uses, finding that the sources of uncertainty are important in determining which conservation payment schemes are likely to be most effective.

Despite the use of economic theory-driven models, individual behaviour and decision making has yet to be extensively considered in the context of land conversion. Individuals' motivations and behavioural factors beyond classical theoretical models are not often included in these analyses. As such, the existing literature may not tell the complete story of how landowners make decisions about land use and land conversion. There is evidence that classical economic theory does not fully describe observed land conversion behaviour in the Prairie Pothole Region (Doidge et al., forthcoming). Despite the potential for higher returns from crop production, conversion rates indicate that farmers are not converting land to capture these increased revenues (i.e. they are not behaving as expected profit maximizers). This suggests that farmers consider more than expected returns when making land conversion decisions.

We propose that anticipated regret may mediate landowners' conversion decisions, such that the potential for regret may make them less willing to convert their land. Intuitively, regret provides a potential explanation for the decision to leave land in grass when conversion to cropland is likely to be more profitable. When farmers are deciding whether or not to convert their land to crop, they may consider how they would feel about their decision ex post. If they anticipate that, in the event of poor growing conditions or a decrease in crop price, they would have received a greater payout from leaving their land in grass and would regret the decision to invest in conversion, keeping their land in grass could be the more desirable choice.

Converting land from grass to crop will likely involve a sunk cost investment by farmers. Depending on the state of the grassland and its previous or current use, farmers may face costs to remove rocks and trees, or existing plant matter. An investment of time, labour, equipment, and other materials may be needed to prepare land for cropping. The existence of this sunk cost may provide an additional barrier to conversion as it may contribute to feelings of regret if the conversion was not profitable. Additionally, as discussed in the next section, there is a body of literature on the regret of action and inaction, with many studies finding that regret is more strongly manifest for decisions that change the status quo rather than those that maintain it (Kahneman and Tversky, 1982; Zeelenberg et al., 2002; Nicolle et al., 2011).

We investigate the role of regret in farmers' land conversion decision in two ways. We first use farmers' stated willingness to pay (WTP) for land conversion from grass use to crop use in hypothetical scenarios under alternative random outcomes. We examine the role of anticipated regret on the amount farmers would pay to convert land and attempt to quantify the impact of regret on farmers' land conversion decisions. We then use data from framed land conversion experiments to investigate the role of regret salience and anticipated regret on rates of land conversion. We also use these experimental data to examine the regret felt from farmers' conversion decisions, specifically looking at the relationship between expected outcomes and regret as well as between regret of action and regret of inaction.

Although many experiments have been conducted to explore regret in decision making, we know of no experiments that have examined the role of regret in economic decisions among a population of experienced professionals. When the decisions of professionals are compared to those of the general population, professionals have been found to behave more rationally than those less experienced in the subject matter (List, 2004; List and Haigh, 2009). Experienced professionals can be expected to have a deeper understanding of the decision context and to be more aware of the consequences of different outcomes. Their responses may therefore differ from those of the general population or from undergraduate students, the population from which economic experiment participants are often drawn.

The results from our WTP scenarios find that, when compared with expected profit maximization, landowners are more disposed to under-conversion from an economic standpoint. Despite this, the functional form of the regret function proposed in this paper is not supported by our data. The results from our experiment, however, suggest that regret may affect farmers' conversion decisions. We observe that farmers are less likely to convert land when regret is made salient. We also find that farmers were more likely to express regret if the condition outcome differed from their expectation, and that more regret was expressed about conversion than non-conversion decisions.

The rest of this paper is organized as follows. Section 2 provides a more comprehensive discussion of past work on regret. Section 3 describes regret theory in more detail and outlines our theoretical framework. Section 4 describes our data collection methods, including experimental procedures. Section 5 presents the empirical strategies used to measure the impact of regret on farmers' land conversion decisions. Section 6 presents and discusses our results, and Section 7 concludes.

3.2. Literature review

The potential for regret exists when agents make a sub-optimal decision, such that another decision would have resulted in a better outcome. Upon learning that her decision was not optimal, the agent may feel regret about her decision ex post. While the study of regret in economics is not widespread, the impact and determinants of regret have been studied extensively in the psychology literature. Research has been conducted on the effect of regret salience on decision making, the impact of anticipated regret on agents' choices, and the effect of decision attributes on feelings of regret.

Regret salience has been found to impact decision making behaviour and processes in economic and psychologic experiments. In sender-responder trust games, senders were found to be willing to send less money when regret was made salient (Kugler et al., 2009). Consistent with the literature on priming and salience, Connolly and Reb (2012) and Connolly et al. (2013) found that regret salience eliminated the decoy effect (a violation of the independence of irrelevant alternatives). Further, regret salience had a moderating effect on decision making, such that subjects who were made to consider regret gathered more information and took longer to make decisions (Reb, 2008). The anticipation of regret has been found to influence agents' decisions, with experiment participants choosing options with smaller potential for regret, even when these decisions were riskier than the foregone choices (Zeelenberg et al., 1996, Zeelenberg and Beattie, 1997).

Research has also explored whether and how regret differs for action and inaction choices, with some studies concluding that feelings of regret are often more intense for decisions that change rather than maintain the status quo. In a study by Kahneman and Tversky (1982), participants were shown two different investment scenarios and asked whether they thought the subject of the scenarios would experience more regret from a loss that resulted from buying (the action decision) or holding (the non-action decision) a particular stock. The vast majority of participants expected that buying would result in stronger feelings of regret, even though both decisions resulted in the same magnitude of loss. Similar findings were reported by Zeelenberg et al. (2002), who asked participants to consider how soccer coaches might feel after losing a game. Participants predicted that the coach who changed his player lineup would feel more regret if he lost the match than the coach who made no changes to his lineup. Nicolle et al. (2011) conducted experiments to test the differential regret of action and inaction of decisions. Study subjects were asked to predict whether a digital ball would fall inside or outside a computerized tennis court, were then provided with feedback about their decision (whether their prediction was right or wrong), and were asked to maintain or change their prediction for the next round (accept or reject the status quo). When study participants made the incorrect decision ex post, feelings of regret were stronger for decisions that rejected, rather than accepted, the status quo (Nicolle et al., 2011).

Not all studies, however, have found that action decisions result in more regret than non-action decisions. When the action is deemed to be justified, the opposite has been found. In Zeelenberg et al. (2002) study participants were asked to consider how soccer coaches might feel about a loss if they had just changed their team lineup (action) or made no lineup change (inaction). The authors found that when the change was felt to be justified (i.e., it occurred after a loss), subjects predicted that inaction would result in stronger feelings of regret than an action decision. The same conclusions were reached by Inman and Zeelenberg (2002), who investigated regret of purchase decisions. If they justified their decision to switch from a trusted brand, agents felt less regret when their decision was sub-optimal than if they felt their decision to switch was not justified.

Agents' expectations about outcomes have been found to impact feelings of regret. Huang and Zeelenberg (2012) conducted a series of experiments that asked subjects to suppose they had invested money in a hypothetical fund. Subjects were then told whether fund returns were greater than or less than the stated expected returns. In both expectation treatments, subjects were told that investment in a different fund would have yielded a higher return. The authors found that when the funds exceeded their expected performance, reported regret was lower than when outcomes fell below expectations, despite the fact that both groups were told they made a sub-optimal decision in their investment choice.

Regret has been used as a motivation for inaction inertia, a behavioural phenomenon of passing up an attractive opportunity (such as a price discount on a wanted item or selling a stock at a profit) after passing on a better opportunity (receiving a larger discount on the item or higher stock price) (Tykocinski et al., 1995; Tykocinski et al., 2004). Theories to explain the role of regret in inaction inertia suggest that people feel regret about the best (missed) opportunity, causing them to value the next opportunity less, or that regret about the missed opportunity will be felt if they seize upon the second opportunity (Tykocinski and Pittman, 1998; Arkes et al., 2002; Sedvalis et al., 2006).

The concept of regret in economic decision making was initially introduced by Savage (1951) who suggested that agents may choose the ex post regret-minimizing action, regardless of the probability that a particular state would occur[18]. Economic regret theory was further developed by Loomes and Sugden (1982), who proposed regret theory to explain violations of some of the axioms of prospect theory observed in Kahneman and Tversky (1979). Loomes and Sugden (1982) developed a theoretical framework to incorporate the role of regret into agent's risky choices. The model involves at least two action choices available to the agent, and at least two possible states of nature that occur with known probabilities. Their modified utility function incorporates the utility she would experience from the choice made in the realized state of nature if she had not made that choice (dubbed her choiceless utility) and also an additional modified utility term. The modified utility term is a function of the difference between the chosen and foregone outcomes, and is such that that, in addition to the choiceless utility, agents experience negative utility if they discover ex post that another decision would have yielded a higher return. This negative utility term is called regret.

[18] This is often referred to as minimax regret, in that agents make the choice that minimize their potential maximum regret, regardless of the probability of experiencing that regret.

In the economic literature, conceptualizations of regret theory often incorporate the probability of experiencing regret. In these anticipated regret frameworks, agents choose based on the magnitude of regret and the probability that they will experience that regret. Braun and Muermann (2004) developed a regret-theoretic expected utility framework that incorporates anticipated regret into an expected utility model, in which a regret function enters into the agent's expected utility function. In their model, agents experience the utility of their choice as well as a disutility of regret from having made a suboptimal choice ex post. Agents seek to maximize their expected utility, which incorporates expectations of regret.

3.3. Conceptual framework

Our conceptual model of regret is framed in a land conversion setting. We suppose that farmers have a plot of land currently in grass. They are faced with the decision of leaving their land in grass or converting it to cropland. Conversion to crop may increase the farmer's returns to the land, but conversion costs must be incurred to receive these economic benefits. Conversion costs that the farmer faces are represented by θ. We model these costs as annualized rather than a one-time conversion cost to match the yearly returns from the land.

We assume two possible states of nature: revenue conditions can be high with probability p or low with probability $1-p$. The annual returns from crop in the high and low states are π_{ch} and π_{cl}. Similarly, respective high and low state returns from grass are π_{gh} and π_{gl}. The upper panel in Table 3.1 summarizes state-conditioned returns. Returns are ranked such that $\pi_{ch} > \pi_{gh} > \pi_{gl} > \pi_{cl}$. This order is chosen to highlight the trade-off between cropping and grass, and also to reflect reality in the production environment we will study[19]. Let $\Delta_h \equiv \pi_{ch} - \pi_{gh} > 0$ and $\Delta_l \equiv \pi_{cl} - \pi_{gl} < 0$ be, respectively, the differences in returns from cropping relative to grass in the high and low states.

If farmers maximize expected profits when making conversion decisions, they will choose the land use with the highest return net of any conversion costs. In our context of land use alternatives, the expected return from the two choices are defined as

$$\begin{aligned} E[\pi_c] &\equiv p\pi_{ch} + (1-p)\pi_{cl}; \\ E[\pi_g] &\equiv p\pi_{gh} + (1-p)\pi_{gl}. \end{aligned} \tag{38}$$

where the *c* and *g* subscripts denote crop and grass. We assign θ_E as the difference between expected returns to crop and grass, such that

$$\theta_E \equiv E[\pi_c] - E[\pi_g] \equiv E[\Delta\pi] \tag{39}$$

In addition to the returns to crop and grass that farmers receive as a result of their decision, the difference between returns to crop and grass in the high and low states may contribute to farmers' feelings of regret. Regret may arise if the wrong decision was made, such that the other choice would have yielded a higher payoff in the state of nature that occurred. Following the economics regret literature, we define the magnitude of regret as the difference between payoffs of the chosen and ex-post optimal action. The lower panel in Table 3.1 provides the state-conditioned regret amount for each choice. A farmer will experience regret if he chooses to convert his land to crop and the low state occurs, since $\Delta_l < 0$. Conversely, he will regret his decision if he chooses not to convert his land and the high condition state is realized,

[19] Commodity prices and resulting returns to crop production have been highly variable in recent years. See Good et al. (2016) for a discussion of crop price variability in the US.

since $-\Delta_h < 0$. If he chooses to convert and the high state occurs or chooses to leave his land in grass and the low state occurs, he will receive the maximum returns in that state of nature and experience no regret in his decision[20].

Two key assumptions are made in our formal model of regret. One is that, if regret is not a factor, growers seek to maximize expected profits. Our model measures anticipated regret as the product of regret magnitude and the probability of the state of nature under which it occurs. Anticipated regret always enters negatively into the agent's payoff function.

3.3.1. Willingness to pay

As denoted by (39), a profit-maximizing farmer should convert her land whenever expected returns to conversion are greater than the costs of doing so. As such, a farmer who seeks to maximize expected profit should be willing to pay any amount up to $\theta_E \equiv p\Delta_h + (1-p)\Delta_l$ to convert her land. If, however, regret enters her payoff function, her willingness to pay for land conversion may differ from the expected returns from doing so.

We borrow from the regret-theoretic expected utility framework of Braun and Muermann (2004), adding an expected regret term in the agent's payoff function. Let $G(z):(-\infty,0] \to (-\infty,0]$ be a regret function, assumed to be increasing in its argument with $G(z)|_{z=0} = 0$ and $\lim_{z\uparrow 0} G(z) = \tau_0 \le 0$. The function is monotonic in the difference between the outcome of returns from the chosen land use and the alternative return given the realised state of

[20] Some discussions of this topic propose a function that takes on a negative value when the difference between the outcomes is negative, such that the agent experiences regret, and a positive value if the difference is positive, such that the agent rejoices in her decision (see Loomes and Sugden, 1982). The model's qualitative results are unaffected by the inclusion of these terms, and so we do not include them in this paper.

nature. Similar to the discontinuity parameter in hyperbolic discounting (Laibson, 1997), the function is possibly discontinuous at 0 in order to capture any existence discontinuity due to the presence of a loss, however small. A farmer will experience regret in her decision if she chooses to convert her land to crop and the low state occurs (i.e., $G(\Delta_l) < 0$). Conversely, she will regret her decision if she chooses to leave her land in grass and the high state is realised, since $G(-\Delta_h) < 0$. If she chooses to convert and the high state occurs or chooses to leave her land in grass and the low state occurs, she will experience no regret in her decision.

Let U_c be the payoff from conversion, including regret. The expected payoff of conversion, $E[U_c]$, is a function of returns from the farmer's land use choice and the random state of nature (the so-called choiceless utility (Loomes and Sugden, 1982)) as well as the anticipated regret of making the wrong choice. Accordingly, the agent's expected payoff of conversion is expressed as

$$E[U_c] = E[\pi_c] + G(1-p, \Delta_l) - \theta \qquad (40)$$

where $E[\pi_c]$ represents the expected returns from conversion in the absence of regret, as defined in equation (38), θ is the annualized cost of conversion, and $G(1-p, \Delta_l)$ is anticipated regret of conversion. The expected payoff of leaving the land in grass can similarly be written as

$$E[U_g] = E[\pi_g] + G(p, -\Delta_h) \qquad (41).$$

Farmers should convert if their expected payoff from cropping exceed that from grass, i.e.,

$$E[U_c] > E[U_g]$$

$$\Rightarrow E[\pi_c] + G(1-p, \Delta_l) - \theta > E[\pi_g] + G(p, -\Delta_h)$$

$$\Rightarrow E[\pi_c] - E[\pi_g] + G(1-p, \Delta_l) - G(p, -\Delta_h) > \theta \text{, or}$$

$$\theta_E + G(1-p, \Delta_l) - G(p, -\Delta_h) > \theta \qquad (42)$$

where $\theta_E = E[\pi_c] - E[\pi_g]$ as defined above and θ are the per-acre conversion costs faced by the farmer.

When anticipated regret plays no role, the conversion decision is one of expected returns only. Farmers should convert grassland to cropland when $\theta_E > \theta$ and not convert when $\theta_E < \theta$, where inertia rules for the case of equality. With regret, we define the cut-off-point for conversion costs as

$$\theta_R = \theta_E - [G(p, -\Delta_h) - G(1-p, \Delta_l)]. \qquad (43)$$

where we label the term in the square brackets as anticipated regret. We define over-conversion as occurring when $\theta_R > \theta_E$ and under-conversion occurs when the opposite is true, i.e., $\theta_R < \theta_E$. When a farmer is willing to pay a higher conversion cost than the expected returns to conversion, more land is converted than the amount predicted by expected profit maximization. Conversely, if the conversion cost a farmer is willing to pay is less than the expected returns to conversion, less land will be converted. Examining the anticipated regret terms, suppose that[21]

$$G(p, z) = \begin{cases} p(\tau_0 + \tau_1 z), \quad \tau_0 \le 0 < \tau_1 & \text{when} \quad z < 0; \\ 0 & \text{when} \quad z \ge 0. \end{cases} \qquad (44)$$

Figure 3.1 illustrates where $\tau_0 < 0$ arises whenever the existence of regret, however small in terms of the monetary payoff, creates a psychological cost. Then anticipated regret is equal to

$$G(p, -\Delta_h) - G(1-p, \Delta_l) = (2p-1)\tau_0 - \tau_1 \theta_E \qquad (45)$$

[21] This is one among many possible functional forms for the regret function.

In that case, $\theta_R > \theta_E$ (willingness to pay for conversion is greater than the expected returns to conversion) if (45) is negative, which is true whenever $\tau_0 = 0$ or whenever both $\tau_0 < 0$ and $p > 0.5$ (given $\theta_E > 0$). However, $\theta_R < \theta_E$ will arise whenever $\theta_E > 0$, $\tau_0 < 0$ and $p < 0.5 + 0.5\theta_E(\tau_1 / \tau_0)$, a number that is less than 0.5.

3.3.2. Decision to convert

In our framed land conversion experiments, participants were asked to decide whether they wanted to convert their land to crop or leave it in grass. As depicted in equation (39), the difference between expected returns to crop and grass is denoted by θ_E. An expected-profit maximizer should convert land from grass to crop whenever the expected returns to conversion are greater than the conversion costs faced by the farmer, or when $E[\Delta\pi] > \theta$. The all-else equal default is to remain in grass.

Different frameworks to model the role that regret plays in agents' decision-making process have been proposed. As first described by Savage (1951), agents may choose the option that has the smallest potential for regret, regardless of the probability that regret will be experienced. Under this minimax regret framework, farmers should choose to convert their land if the regret they would experience in the low state is higher then the regret felt in the high state (i.e. if $|\Delta_l| > |\Delta_h|$). If the opposite is true, farmers should choose to leave their land in grass. Indifference should arise when regret in each state is equal.

Alternatively, a farmer may take the probabilities of the states occurring into account and make the choice that minimizes anticipated regret. Again, letting $G(z):(-\infty,0] \to (-\infty,0]$ be an agent's regret function, assumed to be increasing in its argument. The function is monotonic in the difference between the outcome of returns from the chosen land use and the alternative return given the realised state of nature.

As in section 3.1, we represent the payoff from conversion, including regret, by U_c. The expected payoff of conversion, $E[U_c]$, is a function of returns from the farmer's land use choice and the random state of nature as well as the anticipated regret of making the wrong choice. The agent's expected payoff of conversion is expressed in equation (40) , and the expected payoff from leaving land in grass by (41). Farmers should convert if their expected payoff from cropping exceeds that from grass, which occurs when (42) is satisfied. When anticipated regret plays no role, the conversion decision is one of expected returns only. Farmers should convert grassland to cropland when $\theta_E > \theta$ and not convert when $\theta_E \leq \theta$.

3.4. Data and experimental procedures

Data were collected at focus group meetings of farmers in North and South Dakota in early March of 2016. The purpose of these meetings was to learn directly from farmers what factors they consider when making their actual land use decisions. A total of 76 farmers were convened in four locations, three locations in South Dakota and one in North Dakota. All locations were along the James River Valley, in areas of high grassland to cropland conversion in recent years. Convened farmers resided within a 90-minute drive from their survey location. Farmers were asked to complete surveys about their farm, farming practices, and land conversion in the past ten years (since 2006) in four conversion categories[22].

Data used in this paper were collected by two different methods. The first presented farmers with hypothetical returns to land in grass and crop under normal and bad conditions, and

[22] Conversion of cropland to grassland, conversion of cropland to conservation reserve program (CRP) land, conversion of CRP to cropland, and conversion of grassland to cropland. The last category is the main conversion category of interest in this work.

the probabilities of each occurring, and asked them to state their maximum willingness to pay to convert their land from grass to crop. The second was a land use experiment in which farmers were presented with returns to a plot of land in grass and crop under normal and bad conditions, and the probability of experiencing those conditions. Per-acre costs of converting land from grass to crop were also given. Farmers were asked to decide whether they wished to leave their land in grass or convert it to crop. Both data-gathering methods are described in more detail in the following sub-sections.

3.4.1. Willingness to pay scenarios

Farmers were presented with nine hypothetical land conversion scenarios and asked to report their maximum willingness to pay to convert grassland to crop land. For each of the nine scenarios, participants were asked to suppose that they had a plot of land that was in grass but that could be converted to grow crops. The information given to participants in the hypothetical scenarios included returns to cropping and grass in the two possible states of nature, normal and bad years, and the probabilities that each state of nature would occur. The returns to crop and grass in both states, as well as the probabilities of being in either state, were varied in the scenarios. Farmers were given one of two versions of the WTP scenarios that differed in their returns to crop and grass and the probabilities of normal and bad years occurring, so that there were 18 unique scenarios in total.

The annualized conversion costs presented to farmers in the scenarios ranged from $0 (choosing not to convert given the returns for that scenario) to more than $91 per acre. These annual conversion costs are similar to those reported by farmers in another section of the survey, which averaged \$87 per acre for conversion of grassland to cropland[23]. An option to not convert the land for reasons unrelated to profit was also included to distinguish those who chose not to convert due to economic reasons from those who may have had other reasons for leaving land in grass.

For this analysis, the mid-point of each range of WTP options was used as the farmer's response. If a farmer chose "\$0/I would have to be compensated," his WTP was recorded as zero. Farmers who consistently chose "I would not convert for reasons unrelated to profit" are excluded from this present analysis (8 participants), as are all responses from farmers whose choices were inconsistent, i.e., those who chose "I would not convert for reasons unrelated to profit" in some scenarios and reported positive WTP values in others (25 participants). Three hundred and eighty-one WTP scenario responses from 43 participants meet the criteria above to be included in this analysis. Although participants were presented with nine scenarios, some may have been left unanswered. If, despite this, a participant's responses were consistent he or she was included in the sample.

3.4.2. Experimental procedures

We also conducted a framed land conversion experiment at the focus group meetings. The experiment asked farmers to suppose they had a plot of land currently in grass, but that could be converted to cropland for a yearly conversion cost. Yearly revenue was determined by the chosen land use and stochastic weather and market conditions (hereafter referred to as conditions), which could be either normal or bad. Participants were presented with returns to the

---

[23] Reliable estimates of conversion costs to not exist elsewhere. Obtaining data on conversion costs was another purpose of this survey.

plot of land in grass and crop under normal and bad conditions, the per-acre conversion cost for that round, and the conditional probabilities[24] of normal and bad years occurring. A per-acre cost of converting their plot of land was also given.

The experiment began in year 0, in which the land was in grass. Conditions were revealed to participants individually by a member of the research team. This determined revenue generated by their land for year 0. For the coming year (year 1) participants were asked to make a prediction about the prevailing conditions for that year and decide whether they would leave their land in grass or convert it to cropland at the given conversion cost. Conditions were then revealed individually by a member of the research team, and yearly revenue was determined and recorded on that round's decision sheet. The conversion cost (if incurred) was deducted from the participants' revenue for that year. Conversion costs were incurred for the year of initial conversion from grass to crop and every subsequent year of that round so that conversion decisions were not discouraged in later years of play. If land was left in grass for every year in a particular round, no conversion cost was incurred. Total revenue (the sum of revenue in all preceding years) was also recorded on the decisions sheet. A sample decision sheet is shown in the appendix of this paper. Predictions and decisions were then made for the next round. Play continued for ten years. Two to four rounds of ten years were played in each of the four meeting locations[25].

To investigate the effect of regret salience and the magnitude of regret felt by farmers about their land use decisions, half of the decision sheets given to participants asked them to

---

[24] Farmers were given decision sheets which presented the probabilities of good and bad years if the previous year had been good or bad. The outcome generation process conformed with the Markov property, i.e., only the current state is relevant when forming expectations.

[25] The number of rounds was determined by time constraints at each meeting location. Four rounds were planned for all meetings but were not always completed.

record their feelings about their decisions, after conditions for that year had been revealed and revenue determined. Participants who received the regret version were asked to state how they felt on a scale of 1 to 5, where 1 indicated that participants were unhappy and felt regret about their decision, and 5 indicated they were happy with their land use choice. This version of the experiment was randomized among study participants at each meeting.

3.5. Empirical strategy

3.5.1. Willingness to pay scenarios

3.5.1.1. Expected profit maximization

From equation (44) in our conceptual framework, we let $R_h \equiv -p\Delta_h$ and $R_l \equiv (1-p)\Delta_l$ denote expected regret in the high and low states, with $\theta_E = R_l - R_h$. For empirical purposes we characterize farmers' willingness to pay (WTP) for conversion as

$$\begin{aligned}\tilde{\theta}_R &= \overbrace{R_l - R_h}^{\theta_E} - \overbrace{(p\tau_0 - p\tau_1\Delta_h - (1-p)\tau_0 - (1-p)\tau_1\Delta_l)}^{G(p,-\Delta_h)-G(1-p,\Delta_l)} + \varepsilon \\ &= \tau_0 - 2p\tau_0 + \psi R_l - \psi R_h + \varepsilon;\end{aligned} \tag{46}$$

where $\psi \equiv 1+\tau_1$ and ε is a mean-zero error term. We can test (46) empirically by estimating

$$\tilde{\theta}_R = \alpha_0 + \alpha_1 p + \alpha_2 R_h + \alpha_3 R_l + \varepsilon, \tag{47}$$

where $\alpha_0 \le 0$, $\alpha_1 = -2\alpha_0 \ge 0$, and $\alpha_2 < 0 < \alpha_3 \equiv -\alpha_2 > 1$ are expected.

To test whether farmers behave as expected profit maximizers, we first regress farmers' stated willingness to pay for conversion on expected returns to conversion (θ_E above). The linear regression is specified as

$$WTP_{it} = \beta_0 + \beta_1\theta_{E_{it}} + \beta_2\mathbf{x} + \mu_i + \varepsilon_{it} \tag{48}$$

where $\mathbf{x}$ is a vector of farmer-specific characteristics (including the number of years the participant has been farming, education level, the importance of non-profit factors in actual land decisions, the number of years the farmer expects to continue farming, and whether or not he expects a family member to continue his operation), μ_i is a time-invariant individual-specific term (unobserved) and ε_{it} is a mean-zero idiosyncratic error term. Both fixed and random effects linear regressions were run[26].

The data are such that farmers' WTP values are bounded below by 0 and above by 98. (The highest WTP choice given to farmers was "$91 or more," which was coded as $98 so that all WTP values are in $15 increments.) Because of this, Tobit estimations were also run. Tobit regressions allow a dependent variable that has one or more corner solutions. In our case, the dependent variable for individual i at time t, WTP_{it} is bounded below by $0, so that the observed WTP value, WTP_{it} is equal to $\max(0, WTP_{it}^*)$ where the unobserved $WTP_{it}^* = f(\theta_{E_{it}}) + u_i + \varepsilon_{it}$ is a farmer's unbounded willingness to pay. (It is assumed in this model that $\varepsilon_{it} \mid \theta_E \sim N(0, \sigma^2)$.)

The log likelihood function for an observation WTP_{it}, given that it is abounded below and above by 0 and 98, is

$$\begin{aligned} &\log f(WTP_{it} \mid \theta_{E_{it}}) = 1[WTP_{it} = 0]\log\left[\Phi((0 - \beta\theta_{E_{it}})/\sigma\right] \\ &+1[WTP_{it} = 98]\log\left[\Phi(-(98 - \beta\theta_{E_{it}})/\sigma)\right] \\ &+1[0 < WTP_{it} < 98]\log\left[(1/\sigma)\phi\left(\left(WTP_{it} - \beta\theta_{E_{it}}\right)/\sigma\right)\right] \end{aligned} \tag{49}$$

[26] In the fixed effects estimation all time-invariant variables ($\mathbf{x}$) are dropped, so β_2 cannot be estimated.

where Φ and ϕ denote the standard normal CDF and PDF, respectively (Wooldridge, 2010). Coefficient estimates are obtained by maximizing the log likelihood function with respect to β.

If farmers behave as expected profit maximizers, the amount that they should be willing to pay to convert their land from grass to crop should be equal to the expected returns to conversion. We therefore expect $\beta_1 = 1$, such that there WTP is equal to expected returns. If, however, they do not act as expected profit maximizers, we hypothesize

H1: $\beta_1 \neq 1$ (farmers' willingness to pay for land conversion is not equal to the expected returns of conversion).

3.5.1.2. Anticipated regret

To test for the role of anticipated regret in farmers' decisions, we again ran linear panel and Tobit regressions, modelling farmers' WTP as a function of anticipated regret in the high and low states. The linear regression to test for the role of the anticipated regret terms themselves is specified as

$$WTP_{i,t} = \alpha_0 + \alpha_1 p + \alpha_2 R_h + \alpha_3 R_l + \alpha_4 \mathbf{x} + \mu_i + \varepsilon_{i,t} \tag{50}$$

where $\mathbf{x}$, μ_i, and $\varepsilon_{i,t}$ are as defined in equation (48) above. Robust standard errors clustered at the participant level were calculated in the fixed effects linear regressions.

Because of the frequency of zero responses and the censoring of maximum responses, we also ran Tobit regressions. (The log likelihood function for the Tobit regression is as in equation (49), with θ_E replaced by the vector X_{it}, which denotes the regret terms and any individual-specific variables included in the regression.) Marginal effects from Tobit estimations are reported in the tables below.

Referring again to equation (50), we intend to test the following hypotheses

H2: $\alpha_0 < 0$ (evidence of discontinuity)

H3: $\alpha_2 = -\alpha_3$ (evidence of symmetry across states)

H4: $-\alpha_2 = \alpha_3 = 1$ (evidence of unit response to expected profit)

H5: $\alpha_2 < -1$ (evidence of regret in the high state)

H6: $\alpha_3 > 1$ (evidence of regret in the low state)

3.5.2. Experimental analysis

In our framed land conversion experiment, we investigate the role of regret on farmers' conversion decisions in three ways. We first address whether making regret salient by asking participants to consider the ex post regret they may feel about their decisions has an effect on their land use choices ex ante. We also investigate whether farmers make the regret-minimizing land use decision, looking specifically at whether farmers convert land when the potential for regret is higher when land is in crop than in grass. We then turn to examine the effect of regret from previous decisions on subsequent land use choices. Finally, we consider participants' stated regret by examining the factors that contribute to feelings of regret, including farmers' expectations of the coming year's market and weather conditions. We also estimate differences in regret from conversion decisions and non-conversion decisions to evaluate the differential regret of action and inaction. This section outlines our empirical strategies for these three lines of inquiry.

3.5.2.1. Effect of regret salience and regret magnitude on farmers' land conversion decisions

As stated previously, two versions of the experiment were randomly distributed to participants. Roughly half of participants (31 pf 64) received a version that asked them to report their feelings about their land use decision after each year (regret version); the other half was not asked to assess their decisions (control version). Those who completed the regret version were given sheets asking them to state how they felt after the conditions and their revenue was revealed each year, ranging from 1 (deeply regretted their decision) to 5 (happy with their decision); this column was absent in the control version. (See decision sheet example in the appendix). For the present analysis, this measure of how participants felt was converted into an increasing measure of regret, with 5 indicating that regret was felt and 1 indicating contentedness about the decision.

We first measured the impact of regret salience on farmers land conversion decisions by comparing the rates of conversion between participants who received the regret and control versions of the experiment. We then employed random effects probit panel regressions to estimate the impact of regret salience and the regret-maximizing land use on farmers' conversion decisions. The probability that participant i chooses to convert his land in period t of round j is modelled as

$$p(convert_{ijt}) = \Phi\left(\gamma_0 + \gamma_1 regret_i + \gamma_2 crop\_regret_j + \gamma_3 round_{ij} + \gamma_4 \mathbf{y}_i + \mu_i\right) \qquad (51)$$

where $convert_{itj}$ takes the value 1 if participant i converted his land in year t of round j and 0 otherwise. The function $\Phi(\cdot)$ denotes the standard normal cumulative distribution function (CDF). The variable $regret_i$ is an indicator taking the value 1 if participant i received the regret version of the experiment and 0 if he was in the control group. We created a variable

$crop\_regret_{tj}$ to indicate crop as the regret-maximizing land use[27] in year t of round j, taking the value 1 if the magnitude of regret from crop is greater than that for grass and 0 otherwise (i.e. if the magnitude of regret is higher for land in grass). This variable was included to explore whether farmers made the regret-minimizing land use decision. The dependent variable $round_{ij}$ controls for round-specific variables, including $r_{ch}, r_{cl}, r_{gh}, r_{gl}$, and the probability of a normal year occurring. The vector $\mathbf{y}_i$ contains farmer-specific controls relating to experience (years farming, education, total acres operated, and gross sales on their farms), and μ_i is an unobserved time-invariant individual-specific term.

If regret salience has no impact on farmers' conversion decisions, then making regret salient should not affect how they approach land conversion in these experiments, and we should observe $\gamma_1 = 0$. However, we hypothesize that if regret salience plays a role in farmers' land conversion choices,

H7: $\gamma_1 < 0$, asking farmers to state how they feel about their conversion decisions prompts them to consider these feelings ex ante and makes them less willing to convert their land from grass to crop.

Similarly, if anticipated regret plays no role in farmers' conversion decisions, we should observe $\gamma_2 = 0$. However, if anticipated regret does influence farmers' land use decisions, we expect lower rates of conversion when crop is the regret-maximizing land use and hypothesize

H8: $\gamma_2 < 0$, that farmers are less likely to convert their land when crop is the regret-maximizing choice.

[27] In the vast majority of the experimental scenarios, when crop was the regret-maximizing choice in an absolute sense it was also the anticipated regret-maximizing choice. We are therefore unable to distinguish between minimax regret and anticipated regret.

3.5.2.2. Impact of stated regret on future conversion decisions

Among those who received the regret version of the experiment, random effects probit panel regressions were run to estimate the impact of farmers' stated regret about past decisions on the probability of converting land in subsequent years of the experiment. Our estimation equation is

$$p(convert_{ijt}) = \Phi\left(\eta_0 + \eta_1 stated\_regret_{ijt-1} + \eta_2 round_{ij} + \eta_3 \mathbf{y}_i + \mu_i\right) \tag{52}$$

The dependent and independent variables are the same as those included in (51), with the exception of $regret_{ijt-1}$, which denotes the level of regret farmer i stated about his decision in period $t-1$ of round j. Vectors $round_{ij}$ and $\mathbf{y}_i$ are control for round and farmer specific variables as described above. If stated regret about past decisions has no impact on conversion decisions in the current year, we expect $\eta_1 = 0$. If, instead, past regret makes farmers less likely to convert in the current year, we hypothesize

H9: $\eta_1 < 0$, feelings of regret about past decisions make farmers less likely to convert land in the current period.

3.5.2.3. Factors Influencing Regret

We investigated the determinants of regret, including the effect of expectation and the differential regret of action and inaction, by regressing stated regret on farmers' predictions and conversion decisions. We use random effects probit panel and ordered probit panel regressions to model the regret that participant i expresses about his decision in year t of round j. To capture the effect of expectations on regret, dummy variables were created to indicate whether or not a farmer correctly predicted that year's weather and market conditions. Regret was captured on a scale of 1 to 5, where 1 indicated the farmer was happy with his decision and felt no regret, and 5 indicated that he felt regret. We first convert stated regret into a binary variable, characterizing stated regret of 1 or 2 as feeling no regret, and 3, 4, and 5 as feeling at least some regret. We run random effects probit panel regressions of the probability that participants stated they felt some regret, estimating

$$p\left(some\_regret_{ijt}\right)=\Phi\left(\rho_0+\rho_1 convert_{ijt}+\rho_2 correct_{ijt}+\rho_3 round_{ij}+\rho_4\mathbf{y}_i+\mu_i\right) \qquad (53)$$

where $some\_regret_{ijt}$ is an indicator with the value 1 if participant i stated he felt some regret (i.e. reported his regret as either a 3, 4, or 5) in year t of round j and 0 otherwise, $convert_{ijt}$ indicates that he decided to convert his land in year t, taking the value 1 if he converted his land and 0 if he decided to leave it in its current use. The variable $correct_{ijt}$ takes the value 1 if the farmer's prediction about year t's conditions were correct and 0 otherwise. The vectors $round_{ij}$ and $\mathbf{y}_i$ are as described above.

If regret is felt equally from action and inaction decisions, then we expect $\rho_1=0$, indicating that the type of decision (convert or don't convert) has no impact on the magnitude of regret experienced. Similarly, if expectations have no impact on regret, we expect $\rho_2=0$. Alternatively, we hypothesize

H10: $\rho_1>0$, participants are more likely to feel regret about decisions to convert land (change the status quo) than decisions to leave land in its current use (maintain the status quo), and

H11: $\rho_2<0$**,** participants are less likely to feel regret if their expectations of that year's prevailing conditions are met (i.e. their prediction of the year's conditions were correct).

We then ran ordered random effects probit panel regressions to estimate the effect of farmers' conversion decisions and expectations on their level of stated regret. Allowing $\mathbf{z}_{ijt} \equiv \delta_1 convert_{ij,t-1} + \delta_2 correct_{ijt} + \delta_3 round_{ij} + \delta_4 \mathbf{y}$, we assume that participants' true regret, $regret_{ijt}^* = \mathbf{z}_{ijt} + \varepsilon_i$, is unobserved. If we assume that stated regret, $regret_{ijt}$ is such that $regret_{ijt} = 1$ when $regret_{ijt}^* \le \varphi_1$, $regret_{ijt} = 2$ when $\varphi_1 < regret_{ijt}^* \le \varphi_2$, …, and $regret_{ijt} = 5$ when $regret_{ijt}^* > \varphi_4$ (with $\varphi_1 < ... < \varphi_4$ as unknown cut points to be estimated), the probability that participant i's stated regret is 1 through 5 is modelled as

$$\begin{aligned}
&P\left(regret_{ijt} = 1\right) = P\left(\mathbf{z}_{ijt} \le \varphi_1\right) = \Phi\left(\varphi_1 - \mathbf{z}_{ijt}\right) \\
&P\left(regret_{ijt} = 2\right) = P\left(\varphi_1 \le \mathbf{z}_{ijt} + \le \varphi_2\right) = \Phi\left(\varphi_2 - \mathbf{z}_{ijt}\right) - \Phi\left(\varphi_1 - \mathbf{z}_{ijt}\right) \\
&\vdots \\
&P\left(regret_{ijt} = 5\right) = P\left(\varphi_4 \ge \mathbf{z}_{ijt}\right) = 1 - \Phi\left(\varphi_4 - \mathbf{z}_{ijt}\right),
\end{aligned} \tag{54}$$

(Wooldridge, 2010). Reporting no regret ($regret_{ijt} = 1$) and feeling deep regret ($regret_{ijt} = 5$) are our main categories of interest. To determine the directional impact of the independent variables in the model, coefficients estimated by the model are such that the sign of the effects of the independent variables on $P(regret_{ijt} = 1)$ are opposite sign of the coefficients, and the same sign of the coefficients as their effect on $P(regret_{ijt} = 5)$.

Our hypotheses for the ordered probit model are therefore similar to H3 and H4. We hypothesize that

H12: $\delta_1 > 0$, decisions to convert land make participants less likely to express satisfaction with their decision and more likely to express decision regret, and

H13: $\delta_2 < 0$, participants are more likely to be happy about their decision when their predictions are affirmed, and more likely to express regret about decisions when the condition is not as they predicted.

3.6. Results and discussion

3.6.1. Summary statistics

Summary statistics for experiment participants are presented in Table 3.2. On average, participants were over 50 years of age[28] and had over 37 years farming experience, and expected to continue farming for roughly another 14 years. Most (over 97%) of participants were male, and had completed at least some post secondary education. Participants expected that they would continue to operate their farm for an average of 14.5 more years, and 73% expected that a family member would take over the operation upon their retirement. The average farm size of meeting attendants was over 2,000 acres, with gross farm sales over $250,000 annually. The majority of farmers had experience with conversion on their land: of the 76 participants, 45 had converted land in at least one of the four categories. As can be seen from Table 3.3, over a quarter had converted land from grassland to cropland in the previous ten years.

3.6.2. Willingness to pay scenarios

3.6.2.1. Expected profit maximization

The variables included in the WTP regressions are described in Table 3.4. Table 3.5 presents the results of the regressions that test whether farmers maximize expected profit in these hypothetical land conversion scenarios. The coefficient on the expected returns to conversion approximately 0.3 in all model specifications. That it is consistently less than 1 indicates that farmers are not behaving as expected profit maximizers. A Wald test failed to reject our hypothesis H1 that $\beta_1 \neq 1$. This result shows that in these hypothetical scenarios, farmers are under-converting land with respect to the expected returns. They are willing to pay significantly

[28] Age was captured with a categorical variable (see Table 3.2), so an exact average cannot be calculated.

less for conversion than the expected gains they would receive. These results are consistent with other findings from the surveys completed by farmers about their actual conversion decisions, and with others' finding of land use change elasticities (Lubowski et al., 2008; Rashford et al., 2008). These results imply that farmers land use decisions are affected by non-pecuniary factors, and provide a rationale for investigating alternative explanations for farmers' conversion decisions.

3.6.2.2. Anticipated regret

Table 3.6 compares the difference between participants' WTP and the expected returns to conversion (WTP-θ_E). Our model of anticipated regret predicts that under-conversion should be observed when the probability of being in the normal state, p, is less than 0.5, and over-conversion should occur when $p > 0.5$ (see equation (45)). Contrary to these predictions, over-conversion is observed when the $p < 0.5$, as seen by WTP greater than the expected returns to conversion. As p increases, the difference between expected returns and WTP decreases.

Similarly, the results from the regressions investigating the role of anticipated regret do not support the functional form of the regret function proposed above. As shown in Table 3.7, the estimated constant (α_0 in equation (50)) is significantly greater than zero, suggesting that the expected payoff function with the regret component is not discontinuous in the way specified in this paper. This result is statistically significant at the 1% level in all empirical specifications of the model (fixed and random effects linear regressions, and the random effects Tobit regression). This result does not necessarily provide evidence against the role of regret in land conversion decisions, however, only that the functional form proposed here (specifically the discontinuous functional form) is not supported by these data. Additionally, the coefficient on the probability of being in the normal state is negative in all specifications of the model. This is also contrary to the predictions of the theoretical model, which predicts that as being in the normal state becomes more likely, farmers should be willing to pay more for conversion.

As with estimation of equation (48), unit-response to expected returns to conversion is rejected from estimations of equation (50). The coefficients on the individual regret terms are significantly different from -1 and 1, the model predictions if farmers responded only to changes in expected returns when determining their willingness to pay for conversion. We therefore reject our hypothesis H4.

The signs of the coefficients on the regret terms R_h and R_l are consistent with expectations, and make intuitive sense. As R_h increases, anticipated regret in the high state (of not converting land from grass to crop) decreases. As such, farmers' WTP for conversion should decrease, which is reflected in the negative sign on the coefficient on R_h. Conversely, when R_l increases anticipated regret from conversion decreases. Farmers' WTP should increase, which is seen in our results with the positive coefficient on R_l. Despite the consistency of signs of the coefficients with our model, their magnitudes differ from what was predicted. Further, we reject our hypothesis H3, that $\alpha_2 = -\alpha_3$. Our results do indicate, however, that regret in the low state, R_l, has more of an impact on farmers' WTP than regret in the high state, R_h. As outlined above, regret in the low state is due to converting land and regret in the high state is due to not converting. This provides some support for the assertion that regret is felt more keenly for actions (conversion) than for inaction (non-conversion). We will explore this in more detail in our framed conversion experiment results.

The random effects linear and Tobit regressions allow us to include farmer characteristic variables in the empirical estimation. The only variable that had any impact on farmers' WTP was the importance with which they rated non-monetary factors in their own land use decisions. As shown in Table 3.7, this is negatively associated with the amount that farmers were willing to pay for conversion in the hypothetical land use scenarios. This result held in the linear and non-linear regressions. This suggests that farmers approached these hypothetical scenarios similarly to their actual conversion decisions, such that farmers who rated non-monetary factors (including ecological considerations, preserving land for future generations) with more importance were less likely to convert grassland to cropland on their farms and ranches. All other individual-specific variables (the number of years the participant had been farming, participants' education level, the number of years the participant expected to operate his farm, and whether a family member was expected to take over the operation) had no significant impact on WTP for conversion.

3.6.3. Land conversion experiment

Of 76 farmers who attended our focus group meetings, 64 participated in the experiment. Participants completed two to four rounds of ten years each for a total of 1980 conversion observations. A summary of farmers' conversion decisions is shown in Tables 3.8 and 3.9. In almost 40% of all rounds played, farmers elected to leave their land in grass for the entire round (no conversion was undertaken). In the rounds in which farmers did decide to convert their land, the majority (almost 75%) had only one instance of conversion (the initial conversion of grassland to cropland). As seen in the table, the majority of conversion across all rounds took place in year 1. Although farmers had the option of converting their land back to grass after their initial conversion to cropland, most conversion decisions (78%) were in the grass-to-crop direction.

3.6.3.1. Effect of regret salience and regret-maximizing land use on farmers' land conversion decisions

Approximately half of the experiments used for this analysis (31 of the 64) asked farmers about how they felt about their conversion decisions, while the rest received the control version. Balance tests were conducted for participant characteristics, and no statistically significant difference in covariates was observed between participants in the regret and control groups (see Table 3.10).

Our results suggest that regret salience has a negative impact on farmers' willingness to convert their land. Figure 3.2 and Table 3.11 present the rates of conversion among participants who received the regret and control versions. Farmers in the control treatment converted their land in approximately 15% of all years, while those who received the regret version converted land in only 7% of all years. A Pearson χ^2 test indicates this difference was statistically significant at the 1% level, providing evidence that anticipated regret impacts farmers' conversions decisions. This result indicates a priming effect, in that asking farmers to consider how they might feel after a decision is made impacts how they behave ex ante.

Table 3.12 describes the variables included in the probit regressions. The results of the probit estimation provide stronger evidence of the impact of regret salience, as shown in Table 3.13. We observe that participants who received the regret version of the experiments were roughly four or five percentage points less likely to convert than those who were completed the control version. We therefore fail to reject our hypothesis H7 that $\gamma_1 < 0$, which supports the assertion that priming subjects about regret impacts their willingness to convert land.

Previous work has shown that making regret more salient increases subjects' decision-making time and the information gathered before making a decision (Reb, 2008). In our land conversion experiments, priming participants to consider how they might feel about their decision may have made them contemplate their options more carefully before deciding whether or not to convert their land. In the experiments conducted by Reb (2008), those who were primed to consider feelings of regret took more time to make decisions and gathered more information than those who were not primed to consider regret. This is not to imply that farmers' actual land conversion decisions are made without careful consideration. However, priming farmers to consider how they might feel after making their land conversion decisions may cause them to more carefully contemplate their decision to convert.

When we examine the impact of the regret-maximizing land use, we observe that when the potential for regret was higher for crop than grass, participants were less likely to convert their land (see Table 3.13). We subsequently fail to reject our hypothesis H8 ($\gamma_2 < 0$), that farmers are less likely to convert land when crop is the regret-maximizing land use. This suggests that anticipated regret causes farmers to convert land less frequently. Farmers with higher education levels were less likely to decide to convert their land, and farmers with higher gross farm sales were more likely. All other included farm- and farmer-specific characteristics (years farming, total acres operated, and participants' main source of revenue) had no statistically significant impact on the probability of conversion.

3.6.3.2. Impact of stated regret on future conversion decisions

Probit regressions show that farmers who expressed more regret about their decisions were actually *more* likely to convert land in ensuing years, even when we control for conversion costs and expected returns, as shown in Table 3.14. This effect still exists, but is less strong, when we restrict our dependent variable to conversions to crop only (i.e., when land was in grass the previous year). We therefore reject our hypothesis H9 that $\eta_1 < 0$ (farmers' stated regret about past decisions has a negative impact on subsequent conversion decisions). These results suggest that if farmers feel regret about their decision not to convert their land, and perhaps miss an opportunity for higher returns, they are more likely to convert their land in the current year. Among these farmers, those with higher gross sales were more likely to convert land and convert land to crop, as were those whose main source of revenue was ranching or mixed crop and animal production (as opposed to those who derived the majority of their revenue from crops). Number of years farming, education, and total number of acres operated did not have a statistically significant impact on the probability of conversion or on the probability of conversion of grassland to cropland.

That participants were more likely to convert land from grass to crop if they felt regret about their decision to leave land in grass the previous period is evidence against regret contributing to inaction inertia. Inaction inertia arises when a missed opportunity makes agents less likely to take advantage of a subsequent, although not as attractive, opportunity in the future. As discussed above, regret has been suggested as a reason for this behaviour (Tykocinski et al., 1995; Tykocinski and Pittman, 1998; Arkes et al., 2002; Tykocinski et al., 2004). We see no evidence of this in our framed experimental results.

3.6.3.3. Factors Influencing Regret

As shown in Table 3.15, we observe no statistically significant relationship between the decision to convert land and the probability that participants felt regret about their choice, and consequently reject our hypothesis H10 that $\rho_1 > 0$. That decisions to convert did not make participants more likely to feel regret may be linked to their feeling that the decision was justified given the information available to them. In the literature on regret, previous studies have found that people were less likely to associate regret with decisions that could be justified by past outcomes Zeelenberg et al. (2002). Interestingly, we observe that farmers were more likely to feel regret about decisions to use their land for crop, even when controlling for that year's conversion decision (i.e., if crop was the status quo decision). This may provide insight into farmers' actual land use decisions, which suggests under-conversion of land from an expected profit standpoint.

Consistent with other work on regret, we find that farmers' expectations influence their expressed regret. When farmers predicted conditions correctly, they were less likely to have stated they felt some regret about their land use decision. As shown in Table 3.15, $\rho_2 < 0$ and is statistically significant at the 1% level. We therefore fail to reject our hypothesis H11, that regret is felt more strongly when expectations are different from outcomes. Not surprisingly, farmers felt less regret when their choice was the revenue-maximizing one. None of the farm- or farmer-specific variables included in these regressions (number of years farming, total number of acres operated, education level, gross farm sales, and participants' main revenue source) had a statistically significant impact on the probability that farmers felt regret about their conversion decisions.

Table 3.16 presents the coefficients from the ordered probit regressions. We are mainly interested in the probabilities of expressing no regret (1) and deep regret (5). While the probit model suggested that decision to convert land had no impact on the probability that farmers felt regret about their decisions, the ordered probit model indicates that farmers are less likely to express satisfaction and more likely to express regret about decisions to convert land. This effect is significant at the 1% level, and we therefore fail to reject our hypothesis H12 (that $\delta_1 > 0$).

Consistent with the probit regression, we also fail to reject H13 that $\delta_2 < 0$ and observe that farmers are less likely to express regret about their decisions when they correctly predicted that year's conditions (their expectations are met). This supports findings of other experiments that found that regret is felt more strongly when outcomes fall short of agents' expectations (Huang and Zeelenberg, 2012). We again observe that farmers are more likely to express satisfaction and less likely to express regret when their decision was the revenue maximizing decision, as in the standard probit regressions. None of the farm- or farmer-specific variables included in these regressions (number of years farming, total number of acres operated, education level, gross farm sales, and participants' main revenue source) had a statistically significant impact on the amount of regret farmers felt about their conversion decisions.

3.7. Further discussion and conclusions

While conversion of grassland to cropland in the Prairie Pothole Region is of significant ecological concern due to the loss of avian breeding grounds, the release of stored greenhouse gasses, and increased input use, among others, there is evidence of under-conversion from an economic standpoint. This suggests that farmers are motivated by factors other than the financial returns to their land. In this paper, we proposed regret theory as an alternative to expected profit maximization to explain observed land use patterns, an aspect of land conversion that has not previously been given much consideration. Landowners may anticipate regret if their conversion decision turns out to be a sub-optimal decision ex post; anticipation of this regret may cause them to keep their land in grass rather than convert it for row crop production.

To investigate the potential role of regret in farmers' land use decisions, we conducted focus group meetings with farmers in the region. Meeting participants were experienced agricultural producers in an area with high rates of grass to cropland conversion. We investigated the potential role of regret in farmers' land conversion decisions in two ways. Using hypothetical land conversion scenarios, we asked farmers to report the maximum they would be willing to pay to convert land from grass to crop. We also used framed land conversion experiments, examining various aspects of the relationship between land conversion and regret.

Consistent with observed conversion patterns, farmers' willingness to pay for land conversion indicated under-conversion from an expected profit maximization standpoint. When faced with land conversion scenarios, farmers' WTP was significantly below the expected returns to conversion. However, when we tested our anticipated regret framework with farmers' WTP for conversion, we did not find support for our theoretical model. While the coefficients on the regret terms in the high and low state were of the expected sign, their magnitudes were smaller than the model's predictions. These results cannot conclusively determine whether anticipated regret plays a role in land conversion decisions, but suggest that farmers' regret function is not of the form proposed here.

In contrast, the results from our framed land conversion experiment suggest that regret may influence farmers' land conversion decisions. Several aspects of regret in decision making were considered, including regret salience, the role of anticipated and past regret in land conversion decisions, and the factors that impact expressed regret. We found that those for whom regret was made salient converted their (virtual) land less frequently than those who were not primed to consider regret. Contrary to our hypotheses, the potential magnitude of regret did not impact participants' conversion decision; they were in fact more likely to convert when the potential for regret under crop was greater than under grass. We find that regret about a decision to leave land in grass made farmers more likely to convert their land the following year, providing evidence against the role of regret in inaction inertia.

We also contribute to the literature by investigating the factors that contribute to feelings of regret, including the differential regret of action and inaction decisions. The literature on this topic is mixed, with some studies suggesting that more regret is felt from action decisions than inaction (i.e., decisions that maintain the status quo), while others observe the opposite. Regret among study participants was not higher when they decided to convert their land; however, those whose land was in crop expressed more regret about their decision. Our results also confirm the findings of other studies that examined the relationship between regret and expectations, finding that farmers were more likely to express regret if the condition outcome differed from their expectation (i.e., their prediction was incorrect).

The results presented in this paper suggest that regret may play a role in farmers' land conversion decisions. Those who decided not to convert land on their own operations, despite the potential for higher returns, may be more cognizant of their future feelings about the decision. They may more carefully consider how they might feel about making a choice that turned out to be the wrong one, and take these anticipated feelings into account when making their land use decisions. These results may also point to ways in which those wishing to prevent more land from being converted can appeal to farmers. If landowners are made to consider how they may feel about a decision that turned out to be suboptimal ex post, they may more carefully consider that decision ex ante.

While grassland to cropland conversion in the Prairie Pothole Region is of significant concern, conversion of land in the opposite direction also occurs. Lark et al. (2015) estimated that over 4 million acres of land in the United States were retired from cultivation (including land enrolled into CRP) from 2008-2012. Although the rate of conversion to grass is lower than that out of grass, much of this conversion has been found to occur in the western area of the western corn belt (Wright and Wimberly, 2013). Land converted from grass to cropland is often of marginal quality (Drummond, 2012; Lark et al., 2015). It may therefore make sense to take it out of cultivation if market changes make crop production less economically feasible. While conversion of cropland to grassland may reverse some of the ecological damage caused by its initial conversion, it is likely that a large portion of the damage cannot be undone. Carbon released into the atmosphere by tilling uncultivated land cannot be un-released. Breeding habitat for migratory birds is likely easier to destroy than to re-establish. Moreover, the economic cost of land conversion cannot be recovered.

If farmers and landowners are made to more carefully consider the long-term consequences of converting land for row crop cultivation, and how they may feel about making an economically sub-optimal decision, they may decide not to undertake the initial conversion, thus saving spending resources that they may regret in the future. Agencies wishing to slow rates of land conversion could appeal to farmers' anticipated regret by making them consider how they may feel about their decision to convert if it turns out to be the wrong decision ex post. The results from our framed land conversion experiments suggest that when farmers consider their future feelings about their land conversion decisions, they are less likely to convert land from grass to crop.

While our experiment provides some insight into farmers' land conversion decisions, we recognize some limitations. As in all economic experiments, external validity concerns exist. Farmers' behaviour in an experimental setting may not perfectly correlate with their actual land conversion decisions. Real life land conversion and land use decisions are unlikely to be made on a year-by-year basis as it was portrayed in the experimental results presented here. Rather, farmers will likely make decisions after observing multiple years' outcomes, and make conversion decisions with time horizons longer than one year. However, farmers were asked to approach their experimental land conversion decisions as they would approach the same decisions on their own farm. We hope that framing the experiment in the context of land conversion, with which study participants were familiar, should attenuate some of these concerns.

Regret is posed here as an explanation for the low rates of conversion from an economic standpoint, and its role in causing farmers to delay decisions to convert is tested through experimental methods. However, other possible explanations for under-conversion are possible. As mentioned briefly in the introduction, the issue of land conversion has been explored in the economic literature through the lens of the dynamic nature of the problem, recognising the potential option value of waiting to convert land (such that waiting may reveal information to farmers about land prices that may allow them to make more informed conversion decisions). Previous work has demonstrated that a real options framework has predicted land use patterns at odds with profit maximization (Song et al., 2011, Miao et al., 2015).

APPENDICES

APPENDIX 3A. Tables and figures

Table 3.1. State conditioned returns for crop and grass, and state conditioned regret for choosing crop or choosing grass

Returns	High state (probability p)	Low state (probability $1-p$)
Crop	π_{ch}	π_{cl}
Grass	π_{gh}	π_{gl}
Regret		
Crop	0	$\Delta_l < 0$
Grass	$-\Delta_h < 0$	0

Table 3.2. Summary of meeting participant characteristics

	Mean	Median	Minimum	Maximum
Years farming	37.64	39	10	69
Age [a]	3.54	4	1	5
Gender (% male)	97.30%	-	-	-
Education level [b]	3.21	3	2	5
Expected future years farming	14.5	10	0	60
Expectation that family member will take over (proportion yes)	0.73	1	0	1
Acres operated	2,086.41	1,350	40	21,000
Gross sales [c]	3.07	3	1	5

[a] Age coding: ‘1’ = 29-34, ‘2’ = 35-49, ‘3’ = 50-59, ‘4’ = 60-69, ‘5’ = >70.

[b] Education coding: ‘1’ = primary only, ‘2’ = high school, ‘3’ = some college, ‘4’ = bachelor’s degree attained, ‘5’ = advanced degree attained.

[c] Categorical variables used to capture gross farm sales. 1=Under $99,000, 2=$100,000-$249,000, 3=$250,000-$499,999, 4=$500,000-$999,999, 5=$1,000,000+

Table 3.3. Summary of participants' land conversion history, by conversion category.

	Number of farmers reporting conversion	*Mean acres*	*Median acres*	*Years under original use*	*Likelihood of converting back*
Cropland to grassland	13	114.7	60.0	9	54%
Cropland to CRP	14	77.0	58.5	47	14%
CRP to cropland	15	397.6	141.0	15	30%
Grassland to cropland	21	237.0	77.5	29[a]	26%

[a] If non-native grassland. Six incidences of native grassland conversion were reported.

Table 3.4. Description of variables included in WTP regressions.

Variable	Description
p	Probability of the high state occurring
R_h	Expected regret in the high state ($p\Delta_h$)
R_l	Expected regret in the low state ($(1-p)\Delta_l$)
Years farming	Number of years a farmer has been employed in farming.
Education	Highest level of education that a farmer has completed.
Importance of other factors	How important a farmer rated non-economic factors for land use decisions on his own farm, rated from 1 (not important at all) to 5 (very important).
How long do you expect to operate your farm?	The number of years a farmer expects to operate his farm.
Family member	Whether a farmer expects a family member to operate his farm when he decides to retire.

Table 3.5. Linear and Tobit regressions of participants' WTP on the expected returns to conversion (E[ΔR]) and participant characteristics.

	Linear FE	Linear RE	Tobit
E[ΔR]	0.314***	0.356***	0.370***
Years farming		0.287	0.336
Education		0.833	1.017
Importance of other factors		-8.040***	-8.323**
How long do you expect to operate your farm?		0.425	0.461
Family member		-3.561	-1.728
Constant	22.951***	32.79	26.356
Observations	381	339	339
R-squared	0.073	0.276	-
Number of participants	43	38	38

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

Table 3.6. Mean difference between WTP and expected returns to conversion (θ_E) by probability of the normal state occurring.

	WTP-θ_E	SD
$p = 0.40$	18.57	24.30
$p = 0.50$	10.17	26.36
$p = 0.60$	16.86	28.07
$p = 0.70$	6.63	30.00
$p = 0.80$	-8.17	29.60

Table 3.7. Linear (fixed and random effects) and Tobit regressions to estimate the impact of regret on WTP.

	Linear FE	Linear RE	Tobit
p	-13.60**	-15.97**	-17.55**
	(5.740)	(6.931)	(6.831)
R_h	-0.368***	-0.387***	-0.414***
	(0.0571)	(0.0460)	(0.0497)
R_l	0.414***	0.426***	0.418***
	(0.0912)	(0.0827)	(0.0847)
Years farming	-	0.503	0.534
		(0.419)	(0.372)
Education	-	1.233	1.390
		(4.627)	(4.126)
Importance of non-profit factors	-	-6.870	-7.291*
		(4.661)	(4.072)
How long do you expect to operate your farm?	-	0.693	0.702*
		(0.424)	(0.372)
Do you expect a family member will take over your operation?	-	-5.932	-3.872
		(8.273)	(7.384)
Constant	31.15***	28.46	22.40
	(2.979)	(27.00)	(30.64)

Standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3.8. Mean number of conversions per round, by frequency

Conversions per round	Frequency	Number of rounds (total)
0	39.9%	78
1	43.9%	86
2	5.6%	11
3	6.6%	13
4	2.0%	4
5	1.5%	3
6	0.0%	0
7	0.5%	1

Table 3.9. Mean conversion rate, by year (all rounds)

Year	Mean conversion rate
1	37.4%
2	8.1%
3	3.5%
4	6.6%
5	6.1%
6	7.1%
7	6.6%
8	7.1%
9	5.6%
10	6.1%

Table 3.10. Randomization check of meeting participant characteristics. p-values obtained from two-way t-tests for continuous variables (years farming, farm acres operated), and χ^2 tests for categorical variables (education and gross sales).

	Regret	Control	p-value
Years farming	34.73	38.67	0.23
Education	3.17	3.34	0.70
Farm acres operated	2,341.27	2,015.18	0.67
Gross sales	2.97	3.16	0.67

Table 3.11. Comparison of conversion rates between the regret and control versions of the experiment, for all years and year 1.

Version	N	Mean conversion rate	χ^2 p-value
Control	33	0.115	0.001
Regret	31	0.072	

Table 3.12. Description of variables included in experiment regressions.

Variable	Description
Regret version	Indicator, 1 if participant completed the regret version of the experiment, 0 otherwise
Maximum regret from crop	Indicator, 1 if crop is the regret-maximizing land use, 0 otherwise
p(normal year)	Probability of a normal year if the previous year was normal
Conversion cost	Per-acre conversion cost
Normal year, $t-1$	Indicator, 1 if previous year's conditions were normal, 0 otherwise
Land in grass, $t-1$	Indicator, 1 if land was in grass in previous, 0 otherwise
Stated regret, $t-1$	Participant's stated regret about last year's land use decision
Revenue maximizing decision	Indicator, 1 if participant made the regret maximizing land use choice under the prevailing conditions, 0 otherwise
Normal weather	Indicator, 1 conditions were normal that year, 0 otherwise
Correct prediction	Indicator, 1 if participant made the correct prediction about that year's conditions, 0 otherwise
Convert	Indicator, 1 if participant converted his land that year, 0 otherwise
Crop	Indicator, 1 if participant's land was in crop that year, 0 otherwise
$some\_regret_{ijt}$	Indicator, 1 if participant i felt some regret (regret=3,4, or 5) in year t of round j, 0 otherwise
$state\_regret_{ijt}$	Participant i's stated regret in year t of round j
Farmer-specific controls	Years farming, education, acres operated, and gross sales

Table 3.13. Probit regression results of the impact of regret salience and maximum regret choice on farmers' yearly conversion decisions, convert in year t as dependent variable (Marginal effects reported)

	Conversion decisions	
Regret version	-0.043*	-0.055**
	(0.0244)	(0.025)
Maximum regret from crop (dummy)	-0.077***	-0.078***
	(0.026)	(0.026)
Normal year, $t-1$	-0.003	0.001
	(0.012)	(0.012)
Land in grass, $t-1$	0.091***	0.086***
	(0.019)	(0.019)
Years farming		0.001
		(0.001)
Education		-0.025*
		(0.014)
Farmland acres		0.000
		(0.000)
Gross farm sales		0.018*
		(0.011)
Main source of revenue (Cropping as reference category)		
Ranching		-0.029
		(0.044)
Mixed crop and animal		-0.020
		(0.026)
Round controls	Yes	Yes
Observations	1,980	1,840
Log likelihood	-491.598	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3.14. Probit regression results of probability that farmers convert and convert specifically to crop in year t. Marginal effects reported.

	Convert		Convert to crop	
Stated regret, $t-1$	0.275***	0.359***	0.134	0.182*
	(0.0952)	(0.102)	(0.0919)	(0.101)
Normal conditions, $t-1$	-0.305	-0.104	-0.310	-0.192
	(0.211)	(0.225)	(0.214)	(0.239)
Land in grass, $t-1$	0.967***	0.920***	-[a]	-[a]
	(0.258)	(0.271)		
Years farming		0.00205		0.0172
		(0.0128)		(0.0145)
Education		-0.177		-0.121
		(0.164)		(0.168)
Farmland acres		0.000		0.000
		(0.000)		(0.000)
Gross farm sales		0.443***		0.418***
		(0.142)		(0.149)
Main revenue source				
Ranching		1.153***		1.139***
		(0.411)		(0.439)
Mixed crop and animal		0.648**		0.637**
		(0.276)		(0.292)
Round controls	Yes	Yes	Yes	Yes
Observations	743	662	743	662
Log likelihood				

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

[a] Variable omitted, as land is necessarily in grass for it to be converted to crop.

Table 3.15. Probit regression results of the probability that farmers expressed some regret about their conversion decision that year (reverse regret = 3, 4, or 5), marginal effects reported.

	Some regret	
Convert (dummy)	0.032	0.040
	(0.030)	(0.035)
Crop (dummy)	0.064**	0.067**
	(0.025)	(0.027)
Correct prediction	-0.078***	-0.092***
	(0.028)	(0.031)
Revenue maximizing decision	-0.152***	-0.169***
	(0.0453)	(0.046)
Normal conditions (dummy)	-0.078***	-0.101***
	(0.027)	(0.031)
Years farming		0.003
		(0.004)
Farmland acres		0.000
		(0.000)
Education		0.085
		(0.062)
Gross farm sales		0.074
		(0.050)
Main revenue source		
Ranching		0.030
		(0.205)
Mixed crop and animal		-0.033
		(0.095)
Round controls	Yes	Yes
Observations	911	811
Log likelihood	-278.667	

Table 3.16. Ordered probit regression results of the probability that farmers expressed regret about their conversion decision that year (regression coefficients reported)

	Stated regret	
Convert (dummy)	0.458**	0.543***
	(0.184)	(0.194)
Crop dummy	0.825***	0.808***
	(0.129)	(0.132)
Correct prediction	-0.363***	-0.377***
	(0.111)	(0.116)
Revenue maximizing decision	-1.565***	-1.613***
	(0.112)	(0.119)
Normal conditions (dummy)	-0.895***	-1.080***
	(0.113)	(0.120)
Years farming		0.021
		(0.029)
Farmland acres		0.000
		(0.000)
Education		0.485
		(0.353)
Gross farm sales		0.368
		(0.281)
Main revenue source		
Ranching		0.808
		(1.213)
Mixed crop and animal		0.126
		(0.664)
Round controls	Yes	Yes
Observations	911	811
Log likelihood	-721.531	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Figure 3.1. Graphical representation of a piecewise linear regret function, possibly discontinuous at 0.

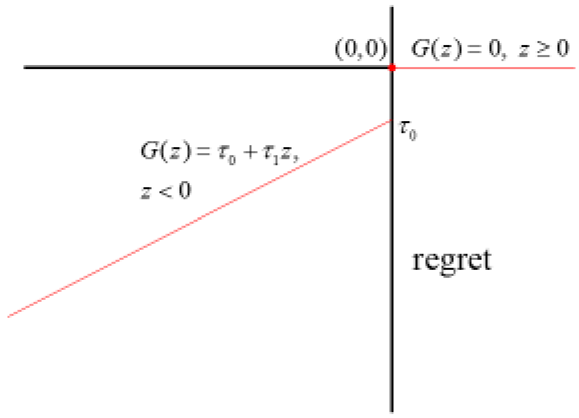

Figure 3.2. Mean conversion rate, by regret treatment.

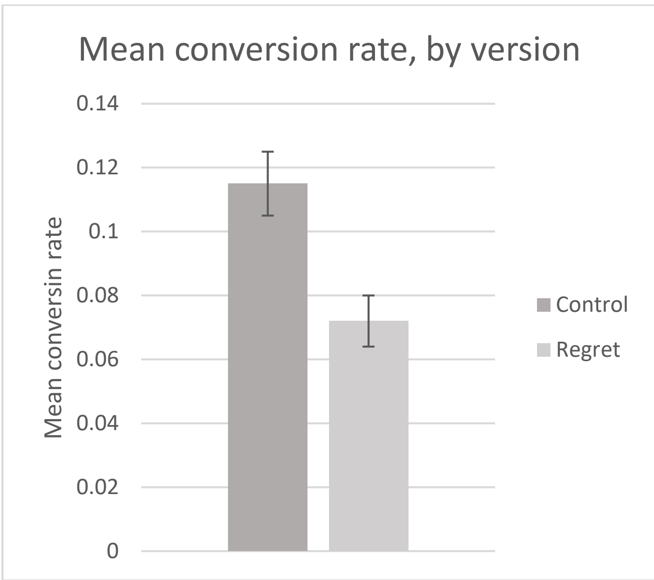

APPENDIX 3B. Land conversion experiment decision sheet

Round 3—decision sheet

Net returns of land uses in normal and bad state

Land use	Weather and market conditions	
	Normal	Bad
Crop	$200/acre	$100/acre
Grass	$130/acre	$110/acre

Probability of conditions in the coming year

If the year before was **NORMAL**	
Normal	0.9
Bad	0.1

If the year before was **BAD**	
Normal	0.4
Bad	0.6

Annual conversion costs this round: $55/acre

Year	Choice		Condition outcome (Normal/ Bad)	Returns this period ($ per acre) **	Total returns ($ per acre)	How do you feel about your decision? 1: deeply regret the decision 5: very happy with the decision	What do you think next year's weather will? (normal/bad)
	Grass	Crop					
0	X					N/A	
1						1 2 3 4 5	
2						1 2 3 4 5	
3						1 2 3 4 5	
4						1 2 3 4 5	
5						1 2 3 4 5	
6						1 2 3 4 5	
7						1 2 3 4 5	
8						1 2 3 4 5	
9						1 2 3 4 5	
10						1 2 3 4 5	

****you must subtract annual conversion costs when you convert from grass to crop in every period after conversion, including the period in which the conversion is made.**

REFERENCES

REFERENCES

Arkes, H.R., Kung, Y.H. and Hutzel, L. 2002. Regret, valuation, and inaction inertia. *Organizational Behavior and Human Decision Processes*, 87(2): 371-385.

Ben-Elia, E., Ishaq, R., & Shiftan, Y. 2013. If only I had taken the other road...: Regret, risk and reinforced learning in informed route-choice. *Transportation*, 40(2): 269-293.

Bailey, J.J. and Kinerson, C., 2005. Regret avoidance and risk tolerance. *Journal of Financial Counseling and Planning*, 16(1): 23-28.

Braun, M., Muermann, A. 2004. The impact of regret on the demand for insurance. *Journal of Risk and Insurance*, 71(4):737-67.

Chomitz, K.M. and Gray, D.A. 1996. Roads, land use, and deforestation: a spatial model applied to Belize. *The World Bank Economic Review*, 10(3): 487-512.

Claassen, R.L., Carriazo, F., Cooper, J.C., Hellerstein, D., Ueda, K. 2011. Grassland to Cropland Conversion in the Northern Plains: The Role of Crop Insurance, Commodity, and Disaster Programs. ERR-120, USDA Economic Research Service.

Cohn, A.S., Gil. J., Berger, T., Berger, H., Toledo Pellegrina, C. 2016. Patterns and processes of pasture to crop conversion in Brazil: Evidence from Mato Grosso State. *Land Use Policy*, 55(September): 108-120.

Connolly, T. and Reb, J. 2012. Regret aversion in reason-based choice. *Theory and Decision*, *73*(1), 35-51.

Connolly, T., Reb, J., and Kausel, E.E. 2013. Regret salience and accountability in the decoy effect. *Judgement and Decision Making*, 8(2): 136-149.

Connolly, T. and Zeelenberg, M., 2002. Regret in decision making. *Current Directions in Psychological Science*, 11(6): 212-216.

Doidge, M, Hennessy, D.A., and Feng, H. Farmers’ Motivations for Land Conversion in the Prairie Pothole Region of North and South Dakota. In L. Knuffman, ed. Proceedings of the 4rd Biennial Conference on the Conservation of America’s Grasslands, Fort Worth, TX, Nov. 14-16, 2017. Forthcoming.

Drummond, M.A., Auch, R.F., Karstensen, K.A., Sayler, K.L., Taylor, J.L., and Loveland, T.R. 2012. Land change variability and human–environment dynamics in the United States Great Plains. *Land Use Policy*, 2: 710-723.

Filiz-Ozbay, E., Ozbay, E.Y. 2007. Auctions with anticipated regret: Theory and experiment. *American Economic Review*, 97(4): 1407-1418.

Fogel, S.O.C. and Berry, T. 2006. The disposition effect and individual investor decisions: the roles of regret and counterfactual alternatives. *The Journal of Behavioral Finance*, 7(2): 107-116.

Geoghegan, J., Villar, S.C., Klepeis, P., Mendoza, P.M., Ogneva-Himmelberger, Y., Chowdhury, R.R., Turner, B.L., and Vance, C. 2001. Modeling tropical deforestation in the southern Yucatan peninsular region: comparing survey and satellite data. *Agriculture, Ecosystems and Environment*, 85(2001): 25-46.

Good, D., Hubbs, T., and Irwin, S. 2017. Are there predictable crop price cycles? *farmdoc daily,* (7):105. Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, June 7, 2017.

Huang, W.H. and Zeelenberg, M., 2012. Investor regret: The role of expectation in comparing what is to what might have been. *Judgment and Decision Making*, 7(4): 441-451.

Inman, J.J. and Zeelenberg, M., 2002. Regret in repeat purchase versus switching decisions: The attenuating role of decision justifiability. *Journal of Consumer Research*, 29(1): 116-128.

Irwin, E.G., and Geoghegan, J. 2001. Theory, data, methods: developing spatially explicit economic models of land use change. *Agriculture, Ecosystems and Environment*, 85(2001): 7-23.

Johnston, C.A. 2014. Agricultural expansion: land use shell game in the U.S. Northern Plains. *Landscape Ecology*, 29(1): 81-95.

Kahneman, D. and Tversky, A. 1979. Prospect theory: An analysis of decision under uncertainty. *Econometrica*, 47(2): 263-293.

Kahneman, D. and Tversky, A. 1982. The psychology of preferences. *Scientific American*, 246(1): 160-173.

Kugler, T., Connolly, T. and Kausel, E.E. 2009. The effect of consequential thinking on trust game behavior. *Journal of Behavioral Decision Making*, 22(2): 101-119.

Laibson, D. 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics,* 112(2): 443-477.

Lark, T.J., Salmon, J.M. and Gibbs, H.K. 2015. Cropland expansion outpaces agricultural and biofuel policies in the United States. *Environmental Research Letters*, 10(4): 044003.

List, J.A., 2004. Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica*, 72(2):615-625.

List, J.A., and Haigh, M.S. 2009. Investment under uncertainty: testing the options model with professional traders. *Review of Economics and Statistics,* 92(4): 974-984.

Li, X., Shou, B. and Qin, Z. 2012. An expected regret minimization portfolio selection model. *European Journal of Operational Research*, 218(2): 484-492.

Loomes, G., Sugden, R. 1982. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal,* 92(4): 805-824.

Loomes, G., Starmer, C., Sugden, R. 1991. Observing violations of transitivity by experimental methods. *Econometrica,* 59(2): 425-439.

Lubowski, R.N., Plantinga, A.J., Stavins, R.N. 2008. What drives land-use change in the United States? A national analysis of landowner decisions. *Land Economics,* 84(4): 529-550.

Malhotra, D., Loewenstein, G., O'Donoghue, T. 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature,* 40(2): 351-401.

Michenaud, S., and Solnik, B. 2008. Applying regret theory to investment choices: Currency hedging decisions. *Journal of International Money and Finance*, 27(5): 677-694.

Miao, R., Hennessy, D.A., and Feng, H. 2015. Sodbusting, crop insurance, and sunk conversion costs. *Land Economics*, 90(4): 601-622.

Nicolle, A., Fleming, S.M., Bach, D.R., Driver, J., and Dolan, R. 2011. A regret-induced status quo bias. *The Journal of Neuroscience,* 31(9): 3320-2217.

Pfaff, A. 1999. What drives deforestation in the Brazilian Amazon? Evidence from satellite and socioeconomic data. *Journal of Environmental Economics and Management,* 37 (1): 26-43.

Rashford, B.S., Walker, J.A., Bastian, C.T. 2011. Economics of grassland conversion to cropland in the Prairie Pothole Region. *Conservation Biology,* 25(2): 276-284.

Ratan, A., Wen, Y. 2016. Does regret matter in first-price auctions? *Economics Letters,* 143(June): 114-117.

Reb, J. 2008. Regret aversion and decision process quality: Effects of regret salience on decision process carefulness. *Organizational Behavior and Human Decision Processes,* 105(2): 169-182.

Savage, L.J. 1951. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253): 55-67.

Sevdalis, N., Harvey, N. and Yip, M. 2006. Regret triggers inaction inertia–but which regret and how? *British Journal of Social Psychology,* 45(4):839-853.

Shah, P., and Ando, A.W. 2016. Permanent and temporary policy incentives for conservation under stochastic returns from competing land uses. *American Journal of Agriculture Economics,* 98(4): 1074-1094.

Song, F., Zhao, J., and Swinton, S.M. 2011. Switching to perennial energy crops under uncertainty and costly reversibility. *American Journal of Agricultural Economics* 3(3): 768-783.

Tversky, A., Thaler, R.H. 1990. Preference reversals. *Journal of Economic Perspectives*, 4(2): 201-211.

Tykocinski, O. E., Pittman, T. S., and Tuttle, E. E. 1995. Inaction inertia: Forgoing future benefits as a result of an initial failure to act. *Journal of Personality and Social Psychology*, 68(5): 793–803.

Tykocinski, O.E. and Pittman, T.S. 1998. The consequences of doing nothing: Inaction inertia as avoidance of anticipated counterfactual regret. *Journal of Personality and Social Psychology*, 75(3): 607-616.

Tykocinski, O., Israel, R. and Pittman, T.S. 2004. Inaction inertia in the stock market. *Journal of Applied Social Psychology*, 34(6): 1166-1175.

United States Department of Agriculture (USDA) National Agricultural Statistics Service. 2017. Farms and Land in Farms 2016 summary. Accessed from https://usda.mannlib.cornell.edu/usda/current/FarmLandIn/FarmLandIn-02-17-2017.pdf

Wooldridge, J.M. 2010. *Econometric Analysis of Cross Section and Panel Data*. MIT press.

Wright, C.K., and Wimberly, M.C. 2013. Recent land use change in the Western Corn Belt threatens grasslands and wetlands. *Proceeding of the National Academy of Science*, 110(10): 4134-4139.

Zeelenberg, M., Beattie, J., Van der Pligt, J., & De Vries, N. K. 1996. Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational Behavior and Human Decision Processes*, 65(2): 148-158.

Zeelenberg, M., & Beattie, J. 1997. Consequences of regret aversion 2: Additional evidence for effects of feedback on decision making. *Organizational Behavior and Human Decision Processes*, 72(1): 63-78.

Zeelenberg, M., & Pieters, R. 2007. A theory of regret regulation 1.0. *Journal of Consumer Psychology*, 17(1): 3-18.

Zeelenberg, M., Van den Bos, K., Van Dijk, E. and Pieters, R. 2002. The inaction effect in the psychology of regret. *Journal of Personality and Social Psychology*, 82(3): 314-327.