Does Selection in Insurance Markets Always Favor Buyers?

Deryugina, Tatyana

6 May 2012

Does Selection in Insurance Markets Always Favor Buyers?

Tatyana Deryugina

University of Illinois at Urbana-Champaign, BIF 4029, 515 E. Gregory Dr, Champaign, IL 61820

Abstract

This paper provides empirical evidence of advantageous selection in insurance markets. By using a novel insurance setting where moral hazard is not a concern, I am able to overcome an important obstacle in most studies of selection: the inability to distinguish moral hazard from selection. In the US market for area yield crop insurance, payouts are based on average county yields. Moreover, area yield insurance is only offered in counties where no farmer is large enough to affect the mean yield. I find that area yield insurance takeup is higher when average yields in the county are higher and show that this effect is not being driven by prices. This suggests that the net selection into these plans thus favors insurance providers, not buyers. One possible mechanism is that providers have better information about aggregate yields. Another is that the desirability of other, non-area yield, insurance options changes, a potentially important but previously overlooked driver of selection.

Keywords: insurance, selection, crop insurance

1. Introduction

The extent to which selection occurs in insurance markets has long been an important empirical question. Traditionally, the focus has been on "adverse" selection, in which higher-risk individuals have more coverage. Adverse selection can lead to large welfare losses and is one of the main economic justifications for mandating insurance coverage. More recently, it has been recognized that, in certain cases, selection may be "advantageous," that is, favoring providers. In this case, lower-risk individuals have more coverage. Advantageous selection can arise when individuals who are more risk averse both buy more comprehensive insurance and take more precautions (de Meza and Webb, 2001; De Donder and Hindriks, 2009). Advantageous selection may also be the result of informational asymmetries that favor providers. When the relationships between characteristics and risk are complex, insurers may have an informational advantage over buyers. Specifically, insurers may have better information about the average risk associated with a particular characteristic. The presence of asymmetric information favoring insurers may undo some or all of the effect of asymmetric information that benefits buyers (Seog, 2009; Villeneuve, 2000, 2005).

Testing for selection in markets such as health and auto insurance is complicated by the presence of moral hazard. These two effects can lead to similar patterns in the data, but they have very different welfare and public policy implications. In this paper, I test for selection in a novel setting where moral hazard is not a concern: area yield plans in the US crop insurance market. Indemnity payments in area yield plans are based on the average yield in a given county and not on individual yields, removing moral hazard concerns. Thus, these plans provide a cleaner test of selection than is possible in other insurance markets.

First, following Finkelstein and Poterba (2004) and Finkelstein and McGarry (2006), I test for selection on public information that affects outcomes but is not reflected in the price. Specifically, I show that yields are autocorrelated but

Email address: deryugin@illinois.edu (Tatyana Deryugina)

area yield premiums do not account for this because of a lag in official data availability. Therefore, recent yields are informative about future yields and should affect farmers' insurance decisions. Thus, the first test consists of checking whether last year's yield affects the total demand for group insurance. Because average county yields should be easily observable, I expect this information to favor farmers.

The second test I develop does not require explicitly specifying the information on which farmers or providers select. In addition to last year's yield, farmers or insurers may have other information about future yields that is not reflected in prices. Even though this information may not be observable to the econometrician, its presence should lead to a correlation between contemporaneous yields and takeup, after controlling for other insurance choice determinants. Three features of group insurance make it possible to attribute any residual correlation between the current yield and takeup to selection rather than to price changes or moral hazard. First, the pricing formula for group insurance plans is known and, as I show, does not reflect contemporaneous yields. Second, because group insurance is typically offered in counties that have many farmers, there is no potential for moral hazard. Finally, the insurance decision is made before crops are planted (and thus before yields are known). Thus, this "reduced-form" test for selection consists of testing whether the number of group insurance policies is significantly correlated with the current yield.

I find some evidence that last year's yield affects takeup of group insurance plans in corn but not soybeans, although estimates of the latter are not very precise. As expected, selection based on last year's yield favors farmers. The second (reduced-form) test indicates that group insurance takeup is higher when average current yields in a county are higher. This suggests that, even though there is some adverse selection, the net selection into area yield plans favors providers, not buyers of insurance.

Finally, I show that last year's yields are also predictive of the total number of insurance policies in the county, including area yield plans and non-area yield plans. This suggests that the desirability of other plans may be changing with yields as well, affecting other options in farmers' choice sets. Changes in the outside option (relative to the plan being considered) is another potential driver of selection that has been previously overlooked. However, because prices in non-group plans are determined using an individual's yield history, which I do not observe, I cannot determine whether the relationship between aggregate takeup of insurance and yield is due to selection, changes in prices, or both.

In addition to being informative about selection more generally, looking at selection in the crop insurance market is important in its own right. Despite a long-run decrease in developed countries' vulnerability to weather shocks, agriculture worldwide remains susceptible to weather fluctuations. According to the Intergovernmental Panel on Climate Change (IPCC), climate change will increase the frequency and intensity of extreme weather events and change their spatial distribution (Meehl et al., 2007; Schneider et al., 2007). This is estimated to be costly to the agricultural sector (Deschenes and Greenstone, 2007). Yield variability and the variance of food prices are likely to increase, at least in the short run. A functioning insurance market may be key to keeping the agricultural sector stable, especially in developing countries. Because of low implementation costs, area yield or weather-based insurance is thought to be the most efficient way to provide crop insurance in many developing countries. This makes studying this particular type of insurance important. For example if there is advantageous selection that is not driven by risk preferences, the gains from having this type of insurance may be eroded.

I contribute to two bodies of literature. First, I contribute to the extensive body of literature testing for adverse selection in various insurance markets. These include health insurance (e.g., Cutler and Reber, 1998; Fang et al., 2008; Einav et al., 2011; Handel, 2011), long-term care insurance (e.g., Sloan and Norton, 1997; Finkelstein and McGarry, 2006; Brown and Finkelstein, 2007), annuities (Finkelstein and Poterba, 2002, 2004; Fong, 2002; McCarthy and Mitchell, 2010), nursing home use (Gruber and Grabowski, 2007), auto insurance (e.g., Chiappori and Salanie, 2000; Cohen, 2005; Chiappori et al., 2006), life insurance (Cawley and Philipson, 1999; He, 2009), and credit cards (Agarwal et al., 2010).[1] The overwhelming takeaway from this literature is that the extent of selection varies significantly and findings from one market may not generalize to another market. Given the importance of the crop insurance market, both currently and in the future, testing for selection in this market is well worth the effort. Moreover, the absence of moral hazard in the plans that I am considering makes this a particularly clean test of selection.

I also contribute to the literature on crop insurance by performing the first test of selection in area yield insurance plans. Although there have been numerous studies of moral hazard and adverse selection in crop insurance, these have typically examined the correlation between farmer characteristics and crop insurance decisions (e.g., Makki and Somwaru, 2001) or focused on a narrow geographical area (Smith and Goodwin, 1996; Horowitz and Lichtenberg, 1993; Coble et al., 1997; Roberts et al., 2006). Several program-wide tests have been performed (e.g., Walters et al., 2007), but to the best of my knowledge, this is the first test to use unpriced public information as the cause of selection and to look at selection in area yield plans.

The rest of the paper is organized as follows. Section 2 describes the empirical setting and the data. In Section 3, I discuss the empirical framework used to test for selection and moral hazard. Section 4 presents the results, and Section 5 discusses and concludes.

2. The Empirical Setting

2.1. The US Crop Insurance Market

Crop insurance plans in the US differ with respect to the insured crop, the metric that determines payment, and the level of the deductible, called the "coverage level." Payments can be determined by (a) individual yield, (b) individual revenue, (c) mean county yield or (d) mean county revenue. Farmers cannot take out multiple insurance plans for the same plot. Within these plan types, farmers can choose from several coverage levels ranging from 50% to 90%.[2] The coverage level specifies the amount by which yield or revenue has to fall (relative to a baseline) before any payment is made. If a farmer chooses a 75% coverage level, for example, he does not receive payments until his yield, revenue, the county yield, or the county revenue (depending on the type of plan) falls more than 25% below the established baseline. The plan type and coverage level largely describe the space of all insurance plans available to farmers planting a given crop.[3] In this section, I describe the area yield plans. Individual yield, individual revenue, and county revenue plans are discussed briefly in Appendix A.

An area yield plan, called a "Group Risk Plan" (GRP) in the US, is based on the extent to which current average county yields deviate from an historic average or "baseline." Baseline county yields, also referred to as the expected yield, are based on at least 30 years of yield history and are trend-adjusted to reflect long-run productivity changes.[4] Per-acre payments in GRP are determined as follows:

$$\text{Pay\_GRP}_{it} = \max\left(0, \left[\text{Yield guarantee}_{ict} - \overline{\text{Yield}}_{ct}\right] \times \text{price election}_i\right) \quad (1)$$

$$\text{Yield guarantee}_{ict} = X_i * \frac{1}{T}\sum_{s=2}^{T+1}\widehat{\text{Yield}}_{c,t-s}$$

where $T > 29$, i indexes the farmer, t indexes the year, and X_i is the chosen coverage level. For group yield plans, per acre payments range from 60% to 90%, in five percentage-point increments. $\widehat{\text{Yield}}_{c,t-s}$ is the de-trended average county

[1] For a broad overview of the evidence for adverse selection in various insurance markets, see Cohen and Siegelman (2010).

[2] The exact range depends on the plan type.

[3] Farmers also have some choices within a plan-coverage-level combination, such as how to combine different plots and how much to get paid in the case of a shortfall. For a more comprehensive overview of the US crop insurance market, see Babcock (2011).

[4] Data on yield histories used to establish county yields are provided by the National Agricultural Statistics Service (NASS).

yield. In practice, the calculation of the expected yield, $\frac{1}{T}\sum_{s=2}^{T+1}\widehat{\text{Yield}}_{c,t-s}$, is more complicated than appears here, as weighted neighboring county yields are used as well. Moreover, a subjective adjustment may be made in some cases. Note that past yields are incorporated into the expected yield (and subsequently into prices) with a lag. This is because NASS releases county yield estimates for the previous year several months after GRP prices and expected yields for the upcoming year have to be set. The presence of this lag potentially creates useful yet unpriced information that insurance providers or farmers can exploit. In the following sections, I estimate both whether lagged yields provide useful information and whether they are subsequently exploited by the buyers or sellers of insurance.

Payments under area yield plans are made if county yields fall far enough below the yield guarantee. Price election is the payment per unit of yield shortfall, chosen by the farmer from a fixed range. For example, a corn farmer with a yield guarantee of 148 bushels per acre, a 75% coverage level, and a \$1.50 price election will get paid \$1.50 for every unit shortfall in yield below 111 bushels per acre. Prices for area yield plans are determined based on parametrically modeling the yield distribution in a particular area (see Skees et al., 1997, for details).

To reduce the level of moral hazard and adverse selection, farmers are required to initiate the purchase of insurance by a certain date, called the "sales closing date." This date varies by county, crop, and year and precedes the earliest allowed planting date.[5] In all cases, insurance decisions are made months before yields for that year are realized. In addition, to ensure that farmers are not "gambling" on yields, they must plant the relevant crop in order to participate in the area yield plan.

2.2. *Insurance Choice*

A farmer will choose the insurance plan that maximizes expected utility. Ignoring risk aversion, the utility of an insurance plan is negatively related to its price and positively related to the expected indemnity payment. The expected indemnity payment, in turn, has two key components: the baseline yield and the actual yield. A higher baseline yield, all else equal, increases the expected payout by making it more likely that the actual yield will be sufficiently below the baseline. A higher actual yield, on the other hand, decreases the payout.

Because there is uncertainty about future yields, informational asymmetries can arise in this market. In particular, there is potential for adverse selection based on private information about likely county yields or on public information that isn't priced into the insurance plan (e.g., forecasted weather conditions). This may lead to adverse or advantageous selection and undermine the functioning of the market.

Both the baseline (or expected) yield and prices are set by the Risk Management Agency every year. I therefore exploit the variation in information about the actual yield to test for selection in and out of area yield insurance plans.

2.3. *Data and Descriptive Statistics*

My data come from three datasets: the Summary of Business Reports, FCI-35 actuarial information, and GRP Final Payment Yields. All are published by the Risk Management Agency (RMA).[6] All data are annual, cover the years 1995-2009, and are broken down by county, crop, insurance plan type, and coverage level. The Summary of Business Reports includes information about the number of insurance policies purchased. The FCI-35 dataset provides information about plan availability and pricing. The GRP Final Payment Yields dataset contains actual yields and expected ("baseline") yields, as computed by the RMA.

Because prices and the overall attractiveness of individual plans vary with individual yield histories while the available data provide only plan-by-coverage-level summaries, testing for selection into individual insurance plans is

[5] In some circumstances, a farmer may purchase insurance after the sales closing date, but these circumstances are limited.

[6] Available from http://www.rma.usda.gov/data/sob/scc/index.html, ftp://ftp.rma.usda.gov/pub/Publications/FCI-35/, and http://www.rma.usda.gov/data/grpfinal/, respectively.

not feasible. Therefore, I focus on area yield insurance plans (GRP), for which all the relevant pricing information is observable. In addition, group plans do not create the risk of moral hazard, making this purely a test of selection. I focus on corn and soybeans because these are the most commonly grown crops and thus have the largest number of relevant observations. I perform the analysis separately for each crop.

Table 1 shows the summary statistics for the key variables used in the analysis, including yields, number of group insurance policies, and the total number of insurance policies. Columns 1 and 2 show the mean and standard deviations for corn and soybeans, respectively. The average number of insurance policies is similar for corn (169 policies) and soybeans (179 policies). GRPs are available in 1,076 counties for corn and 961 counties for soybeans at least once during the time period of interest. GRP plans represent, on average, about 2-3% of all insurance policies. Thus, they are fairly unpopular. A likely explanation for this is that individual plans offer more individually tailored risk protection and generally have higher indemnity payment-to-premium ratios.

To motivate the formal analysis, I begin by illustrating the raw relationships between actual, expected, and last year's yields. Figure 1 shows the relationship between actual and expected yields for corn and soybeans. Although there is a substantial amount of variance, the two are strongly and positively related for both crops. However, a fitted quadratic relationship lies above the 45 degree line for corn, indicating that expected yields are generally below actual yields. For soybeans, the fitted line is somewhat flatter than the 45 degree line, but does cross it. Overall, the expected yield is most biased at high and low yields, but the magnitude does not appear to be large. I later re-examine this relationship taking fixed effects into account.

Figure 2 shows the relationship between actual yield and last year's yield for the two crops. Again, the two are strongly and positively related. The plots for soybeans and corn look very similar. A quadratic fit line indicates that low (high) yields last year are indicative of higher (lower) yields this year. This is true for both crops and suggests that there is some mean reversion in yields. If this is not accounted for in the price, there is potential for exploiting this information.

3. Empirical Framework

3.1. Selection on last year's yield

To test for the presence of adverse selection, I follow Finkelstein and McGarry (2006). Specifically, I estimate the following equations:

$$Log(Yield_{ct}) = \beta_1 Log(Yield_{c,t-1}) + \beta_2 Log(E_{t-1}[Yield_{ct}]) + \alpha_t + \alpha_c + \varepsilon_{ct} \quad (2)$$

$$Policies_{ct} = \gamma_1 Log(Yield_{c,t-1}) + \gamma_2 Log(E_{t-1}[Yield_{ct}]) + \alpha_t + \alpha_c + \varepsilon_{ct} \quad (3)$$

where c is the county and t is the year. $E_{t-1}[Yield_{ct}]$ is the expected corn or soybean yield in county c and year t, as reported by the RMA prior to the start of the growing season. $Yield_{c,t-1}$ is the yield in the previous year, also reported by RMA. $Policies_{ct}$ is the number of corn GRP policies held. α_t and α_c are year and county fixed effects, respectively. The specification for equation (2) is OLS, while the specification for equation (3) is negative binomial. Standard errors are clustered by state-year in this and all following specifications.[7]

Equation (2) tests whether last year's yield is predictive of current yield. I control for the expected yield because it is used to set the price of the insurance policies as well as to determine the yield guarantee. As I explain in Section 2, last year's yield is not used to determine either the expected yield or the price of GRPs. I later confirm that last year's

[7] Clustering by county generally decreases the standard errors.

yield is not reflected in prices empirically. Therefore, if β_1 is significantly different from 0, then there is potential for selection into the GRP plans.

Equation (3) tests for the presence of such adverse selection. In particular, I estimate whether last year's yield affects the number of policies chosen, conditional on the expected yield. The coefficient of interest is thus γ_1. Depending on the relationship between the current yield and last year's yield, it may be positive or negative. All else equal, farmers have an incentive to take out insurance if they know that yields will be lower than the insurance company expects it to be. Thus, if β_1 is positive (negative), γ_1 should be negative (positive) when adverse selection is present.

3.2. Reduced-form test of selection

There is another test of selection that can be performed in this setting because of three specific features of area yield insurance in the US. First, there is no moral hazard in these plans because they are based on the average county yield and are available only in counties with many farms. Second, the insurance choice in a given year is made months *before* yields in that year are realized. Finally, as I show in Appendix B, contemporaneous yields and prices are not correlated.

Suppose there is some information accessible to farmers or providers (e.g., weather forecasts or changes in soil quality) that is not reflected in the price but is predictive of yield in the upcoming season. If buyers or providers act on this information, there should be a significant relationship between current yields and the number of policies, even though current yields are not known when the insurance decisions are made. To test for this, I estimate the relationship between the current yield and the number of GRP policies, controlling for the expected yield:

$$Policies_{ct} = \theta_1 Log(Yield_{ct}) + \theta_2 Log(E_{t-1}[Yield_{ct}]) + \alpha_t + \alpha_c + \varepsilon_{ct} \tag{4}$$

Instead of looking at selection on past observable information, this specification tests for selection on future yield realizations, controlling for the expected yield. Because farmers must purchase insurance long before the current yield is known, θ_1 implicitly captures the effect of other (unpriced) information that farmers or providers may have about yield realizations without having to measure this information explicitly.

3.3. Prices

Finally, I test whether or not prices (as reported in FCI-35 actuarial files) are affected directly by last year's yield or the current yield:

$$Log\left(P_{gct}\right) = \beta_1 Log(Yield_{c,t-1}) + \beta_2 Log(E_{t-1}[Yield_{ct}]) + \alpha_{gt} + \alpha_{gc} + \varepsilon_{gct} \tag{5}$$

$$Log\left(P_{gct}\right) = \delta_1 Log(Yield_{ct}) + \delta_2 Log(E_{t-1}[Yield_{ct}]) + \alpha_{gt} + \alpha_{gc} + \varepsilon_{gct} \tag{6}$$

g = coverage level; t = year; c = county

I perform this test at the coverage level, a more disaggregated level than is involved in the adverse selection test, for the following reason: If the relative proportion of people at each coverage level changes because of adverse selection, the average price across coverage levels might be statistically but not causally related to last year's yields. I use the unsubsidized prices set by the insurance providers. Because premium subsidy rates vary only across years and not across counties, this does not affect the results. The results of this regression are shown in Tables A1 and A2 and confirm that last year's yields and current yields are unrelated to GRP prices.

4. Results

Table 2 shows the relationship between last year's yield and the current yield for each crop, controlling for expected yield, county fixed effects, and year fixed effects. Columns 1 and 2 have the log of the current yield as the dependent variable, while Columns 3 and 4 have the level of the current yield as the dependent variable. The full sample includes all counties in which GRP plans are available at some point during 1995-2009. In Columns 2 and 4, I restrict the sample to counties in which GRP policies are taken up at all during the sample period. The results are similar in both samples. Panel A shows these results for corn. There is no significant relationship between this year's yield and last year's yield in levels and a very weak one between the levels of this year's yield and the expected yield. For soybeans (Panel B), a one-unit increase in yields in a given year is associated with a 0.08-0.09-unit decrease in yields the next year, and this relationship is statistically significant. The relationship between the logs of the yield and the lagged yield is strong and significant for both crops.

According to Columns 1 and 2, a 1% increase in corn yields is associated with a $0.09-0.12\%$ decrease in yields in the following year, holding expected yield constant. Thus, yields appear to exhibit some mean reversion. This process is not necessarily random: The lower yields following higher yields may be due to patterns of nutrient depletion in the soil or of crop rotation. The expected yield is also negatively correlated with the actual yield: A 1% increase in the expected yield is associated with a $0.32-0.38\%$ reduction in the actual yield. This is conditional on county and year fixed effects, as well as on last year's yield; regressing current corn yield on the expected yield with no controls results in a partial correlation coefficient of 0.97 in levels and 0.53 in logs. The results for soybeans are very similar.

These results show that current yield is predictive of future yield. For example, if last year's yield was average, the current yield is expected to be lower than average. Holding price and expected yields constant, this should increase the demand for GRP insurance because the likelihood of a payout is higher. As I later show, this relationship is not taken into account in prices, which means that adverse selection may arise. Specifically, if farmers know that this year's yield will be lower (higher) than expected, they should increase (decrease) their demand for insurance. Because last year's and current yields are negatively correlated, the relationship between the number of GRP policies and last year's yield should be positive if selection favors farmers and negative if it favors providers.

Table 3 shows the relationship between (a) the log of last year's yield and the number of GRP policies and (b) the log of this year's yield and the number of GRP policies, controlling for the log of expected yield, county fixed effects, and time fixed effects. The regression specification is a negative binomial. There is indeed some evidence of adverse selection for corn - the relationship between last year's corn yields and current takeup is positive and marginally significant. There is no significant relationship between last year's yield and the number of soybean area yield policies, although the large standard errors make the estimate imprecise.

Table 3 also shows that current yields are positively correlated with insurance takeup for both corn and soybeans. This suggests that, even though some adverse selection may be occurring, the net selection into GRP plans acts in favor of insurers: In years when yields will be high (controlling for expected yield) more corn and soybean farmers will take out county yield insurance.[8]

It is possible that equation (3) fails to control adequately for other changes that may be correlated with last year's yields. In particular, last year's yields may affect the desirability of individual yield plans or individual revenue plans, either through prices or payout probabilities. This does not matter for the validity of the test of selection into area yield plans. However, in order to understand the behavior of market participants and whether selection might occur on

[8]Appendix B shows the results of an OLS regression of the log of GRP policies on current and past yields, controlling for expected yield and fixed effects. In general, the results for corn are insignificant and those for soybeans are marginally significant. However, due to the count nature of the data, a negative binomial specification is more appropriate.

a more aggregate level, it is useful to look at the aggregate relationship between yields and takeup.

To see whether yield fluctuations are related to aggregate demand for insurance, I regress total insurance takeup on last year's yields and current yields, controlling for the expected yield. The results are shown in Table 4. Columns 1-3 show the results of an OLS regression of log policies on yields, while Columns 4-6 show the corresponding results using a negative binomial specification. A 1% increase in last year's yield is associated with a 0.11% decrease in the total number of insurance policies for corn and a 0.08% decrease in policies for soybeans, while the relationships between current yields and takeup are not significant. The same pattern holds for past yields and current yields in the negative binomial specifications in Columns 3-6, although the coefficients are generally less significant.

5. Discussion and Conclusion

In this paper, I estimate the extent of selection in the US area yield insurance. I implement two tests of selection. First, I use last year's yield (which is predictive of the current yield but does not affect the price of a plan) as a potential source of unpriced information on which farmers or providers could act. I find some evidence for adverse selection based on this information for corn insurance plans, but not for soybeans, although the latter estimates are imprecise.

I then employ a more general test by estimating whether the current yield (which is realized after an insurance purchase decision is made) is predictive of area yield insurance takeup. This is a valid test for selection in settings where the insurance decision clearly precedes the outcome of interest and where moral hazard concerns are absent. This reduced-form test indicates that area yield insurance takeup is higher when average current yields are higher. This suggests that the net selection into area yield plans favors providers, not buyers of insurance.

The design of these plans precludes moral hazard, which enables me to attribute all estimated effects to selection. Because the policy implications for moral hazard and selection are different, it is important to be able to distinguish between the two effects. Moreover, farmers who adopt area yield insurance cannot exploit information about their own idiosyncratic risk because that risk is irrelevant for payouts. This enables me to also cleanly test for selection on aggregate risk. In other markets, selection on aggregate risk may be confounded by countervailing selection on idiosyncratic risk or moral hazard, making the selection on aggregate unpriced information difficult to tease out.

In most insurance markets in which adverse selection is a potential problem, unobserved or unpriced information is assumed to be better known by individuals. In the case of county yields, however, it is plausible that providers have superior information. In particular, they may observe detailed current and past yield information for many farmers in a county (through information needed to buy individual insurance plans, which are very popular) and are thus able to form better aggregate yield forecasts than individual farmers are. However, providers of crop insurance are prohibited from competing on prices. Thus, much of the competition focuses on who to sign up for a particular insurance plan and when. Because crop insurance companies cannot compete on prices and many farmers buy insurance through agents, agent-based competition might be very important in this market. Well-informed agents may be able to convince farmers to choose one plan over another or target particular counties in years when yields there are likely to be high.

I also find that current and past yields are significant predictors of total insurance demand, including individual plans. This suggests that the desirability of non-group insurance plans rises and falls with yields as well, suggesting another mechanism through which the observed selection is occurring. However, because prices in these plans are determined using individual yields, which I do not observe, I cannot determine whether the relationship between aggregate takeup of insurance and yields is due to selection or changes in prices. Therefore, I cannot determine the exact mechanism of selection with the current data. However, this is an important avenue of future research. If providers of area yield or weather-based insurance are able to better predict outcomes than the individuals they insure are and are able to use non-price mechanisms to increase takeup in years when insurance is least necessary, this lowers the welfare gains from area yield or weather-based insurance.

Plans similar to GRP have been advocated and implemented in developing nations because the enforcement costs are much lower and the risk of moral hazard is completely eliminated, as long as the covered area is large enough so that no farmer is able to affect the average county yield on his own. A variation on this approach that is common and relevant is the use of rainfall as the determinant of payouts. One clear drawback of area yield or weather-based plans is that they typically do not provide as much protection as plans based on individual yields provide. Another drawback, as I have shown, is that selection remains a concern in these plans, potentially undermining their ability to provide cheap insurance to farmers.

The generalizability of my results is subject to a key caveat: Area yield insurance or GRP makes up a small fraction of insurance plans in the US. It is possible that selection would operate differently in a setting where only GRPs were available. Although this is beyond the scope of this paper, it would be desirable to replicate this test using insurance data from a developing country, where similar plans are often the only formal insurance available.

6. Acknowledgements

I am grateful to Amy Finkelstein and Michael Greenstone for their invaluable feedback. I thank Jason Abaluck, Jeff Brown, and Julian Reif for helpful discussions and suggestions. I thank Robert Dismukes and Joseph Glauber for helping me to understand the crop insurance market.

References

Agarwal, S., Chomsisengphet, S., Liu, C., 2010. The Importance of Adverse Selection in the Credit Card Market: Evidence from Randomized Trials of Credit Card Solicitations. Journal of Money, Credit and Banking 42, 743–754.

Babcock, B., 2011. The Politics and Economics of the US Crop Insurance Program. Working paper.

Brown, J., Finkelstein, A., 2007. Why is the Market for Long-Term Care Insurance so Small? Journal of Public Economics 91, 1967–1991.

Cawley, J., Philipson, T., 1999. An Empirical Examination of Information Barriers to Trade in Insurance. The American Economic Review 89, 827–846.

Chiappori, P.A., Jullien, B., Salanie, B., Salanie, F., 2006. Asymmetric Information in Insurance: General Testable Implications. The RAND Journal of Economics 37, 783–798.

Chiappori, P.A., Salanie, B., 2000. Testing for Asymmetric Information in Insurance Markets. Journal of Political Economy 108, 56–78.

Coble, K.H., Knight, T.O., Pope, R.D., Williams, J.R., 1997. An Expected-Indemnity Approach to the Measurement of Moral Hazard in Crop Insurance. American Journal of Agricultural Economics 79, 216–226.

Cohen, A., 2005. Asymmetric Information and Learning: Evidence from the Automobile Insurance Market. The Review of Economics and Statistics 87, 197–207.

Cohen, A., Siegelman, P., 2010. Testing for Adverse Selection in Insurance Markets. Journal of Risk and Insurance 77, 39–84.

Cutler, D., Reber, S.J., 1998. Paying for Health Insurance: The Trade-off between Competition and Adverse Selection. Quarterly Journal of Economics 113, 433–466.

De Donder, P., Hindriks, J., 2009. Adverse Selection, Moral Hazard and Propitious Selection. Journal of Risk and Uncertainty 38, 73–86.

Deschenes, O., Greenstone, M., 2007. The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather. American Economic Review 97, 354–385.

Einav, L., Finkelstein, A., Ryan, S.P., Schrimpf, P., Cullen, M.R., 2011. Selection on Moral Hazard in Health Insurance. NBER Working Paper 16969.

Fang, H., Keane, M.P., Silverman, D., 2008. Sources of Advantageous Selection: Evidence from the Medigap Insurance Market. Journal of Political Economy 116, 303–350.

Finkelstein, A., McGarry, K., 2006. Multiple Dimensions of Private Information: Evidence From the Long-term Care Insurance Market. American Economic Review 96, 938–958.

Finkelstein, A., Poterba, J., 2002. Selection Effects in the United Kingdom Individual Annuities Market. The Economic Journal 112, 28–50.

Finkelstein, A., Poterba, J., 2004. Adverse Selection in Insurance Markets: Policyholder Evidence from the UK Annuity Market. Journal of Political Economy 112, 183–208.

Fong, W.M., 2002. On the Cost of Adverse Selection in Individual Annuity Markets: Evidence From Singapore. Journal of Risk and Insurance 69, 193–208.

Gruber, J., Grabowski, D.C., 2007. Moral Hazard in Nursing Home Use. Journal of Health Economics 26, 560–577.

Handel, B.R., 2011. Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts. Working Paper.

He, D., 2009. The Life Insurance Market: Asymmetric Information Revisited. Journal of Public Economics 93, 1090–1097.

Horowitz, J.K., Lichtenberg, E., 1993. Insurance, Moral Hazard, and Chemical Use in Agriculture. American Journal of Agricultural Economics 75, 926–935.

Ker, A., McGowan, P., 2000. Weather-Based Adverse Selection and the U.S. Crop Insurance Program: The Private Insurance Company Perspective. Journal of Agricultural and Resource Economics 25, 386–410.

Makki, S.S., Somwaru, A., 2001. Evidence of Adverse Selection in Crop Insurance Markets. The Journal of Risk and Insurance 68, 685–708.

McCarthy, D., Mitchell, O.S., 2010. International Adverse Selection in Life Insurance and Annuities. International Studies in Population 8, 119–135.

Meehl, G., Stocker, T., Collins, W., Friedlingstein, P., Gaye, A., Gregory, J., Kitoh, A., Knutti, R., Murphy, J., Noda, A., Raper, S., Watterson, I., Weaver, A., Zhao, Z., 2007. Global Climate Projections, in: Solomon, S., Qin, D., Manning, M., Chen, Z., Marquis, M., Averyt, K., Tignor, M., Miller, H. (Eds.), Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

de Meza, D., Webb, D.C., 2001. Advantageous Selection in Insurance Markets. The RAND Journal of Economics 32, 249–262.

Roberts, M.J., Key, N., O'Donoghue, E., 2006. Estimating the Extent of Moral Hazard in Crop Insurance Using Administrative Data. Applied Economic Perspectives and Policy 28, 381–390.

Schneider, S.H., Semenov, S., Patwardhan, A., Burton, I., Magadza, C., Oppenheimer, M., Pittock, A., Rahman, A., Smith, J., Suarez, A., Yamin, F., 2007. Assessing Key Vulnerabilities and the Risk from Climate Change, in: Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, editor=M.L. Parry and O.F. Canziani and J.P. Palutikof and P.J. van der Linden and C.E. Hanson, publisher=Cambridge University Press, address=Cambridge, UK.

Seog, S.H., 2009. Insurance Markets with Differential Information. The Journal of Risk and Insurance 76, 279–294.

Skees, J.R., Black, J.R., Barnett, B.J., 1997. Designing and Rating an Area Yield Crop Insurance Contract. American Journal of Agricultural Economics 79, 430–438.

Sloan, F.A., Norton, E.C., 1997. Adverse Selection, Bequests, Crowding Out, and Private Demand for Insurance: Evidence from the Long-term Care Insurance Market. Journal of Risk and Uncertainty 15, 201–219.

Smith, V.H., Goodwin, B.K., 1996. Crop Insurance, Moral Hazard, and Agricultural Chemical Use. American Journal of Agricultural Economics 78, 428–438.

Villeneuve, B., 2000. The Consequences for a Monopolistic Insurer of Evaluating Risk Better Than Customers: The Adverse Selection Hypothesis Reversed. The Geneva Papers on Risk and Insurance Theory 25, 65–79.

Villeneuve, B., 2005. Competition Between Insurers With Superior Information. European Economic Review 49, 321 – 340.

Walters, C.G., Shumway, C.R., Chouinard, H.H., Wandschneider, P.R., 2007. The Effects of Transitional Yields on Adverse Selection in Crop Insurance Working paper.

7. Appendix A - Insurance Payments in Individual Plans

Because farmer productivity varies, individual baseline yields are necessary for individual-level plans to correctly determine the basis on which payment should be made. The baseline yield for an individual yield plan is established by averaging a farmer's historic certified yields (four consecutive years is the minimum and 10 years is the maximum). Once 10 years of continuous yield history is available, the baseline yield becomes a 10-year moving average, updated every year. If less than four years of continuous yield history is available, average county yields are used in place of an individual yield until the farmer builds up an adequate yield history. If county yields are used to calculate the baseline yield, they are discounted by 15 – 35%, depending on how many years of actual yield history are available.

Per-acre payments under an individual yield plan (called Actual Production History or APH) are determined by the following formula:

$$\text{Pay\_APH}_{it} = \max(0, [\text{Yield guarantee}_{it} - \text{Yield}_{it}] \times \text{price election}_i) \tag{7}$$

$$\text{Yield guarantee}_{it} = X_i * \frac{1}{10}\sum_{s=1}^{10} \text{Yield}_{i,t-s}$$

where i indexes the farmer, t indexes the year, and X_i is the chosen coverage level. For individual yield plans, it ranges from 50% to 85%, in increments of five percentage points. Price election is the payment per unit of yield shortfall, chosen by the farmer from a range set by the Federal Crop Insurance Corporation (FCIC) or the Risk Management Agency (RMA). For example, a corn farmer with a yield guarantee of 148 bushels per acre, a 75% coverage level, and a $1.50 price election will be paid $1.5 for every unit shortfall in yield below 111 bushels per acre.

The baseline for individual revenue insurance is also based on the farmer's 4 – 10 year yield history but takes prices into account as well. The baseline revenue is the average of the individual's historic yields multiplied by the Chicago Board of Trade pre-growing season futures prices for that crop.[9] The actual revenue is calculated using the 1-month futures price near the harvest time for that crop (called the "harvest futures price"):

$$\text{Pay\_RA}_{it} = \max(0, \text{Revenue guarantee}_{it} - \text{Revenue}_{it}) \tag{8}$$

$$\begin{aligned} \text{Revenue guarantee}_{it} &= X_i * P(pre-season)_t * \frac{1}{10}\sum_{s=1}^{10} \text{Yield}_{i,t-s} \\ \text{Revenue}_{it} &= \text{Yield}_{it} * P(harvest)_t \end{aligned}$$

[9]The exact month varies by crop.

In some plans, farmers have the option to have the revenue guarantee based on the *maximum* of the February and the harvest futures prices. This decreases the probability of a claim as well as the price of the insurance plan.

There are four types of individual revenue plans: Crop Revenue Coverage (CRC), Income Protection (IP), Indexed Income Protection (IIP), and Revenue Assurance (RA). The characteristics of all individual revenue plans are very similar. The only difference between CRC and RA plans is that the former limit the amount of payment that is made in case of a loss. The IP plan is similar to the CRC and RA plans, but with this plan it is not possible to base the revenue guarantee on the harvest futures price. In addition, the IP requires that all cropland in a county that belongs to the same entity and grows a particular crop be insured together. This is called an "enterprise unit." CRC and RA allow the cropland to be divided into more basic units and insured separately. Starting in 2011, the four revenue protection plans have been combined into two plans to eliminate redundancies.

The IIP plan is a variation of the IP plan that is based on individual yield histories *relative* to county yields. The baseline yield is established by subtracting the average historic difference between individual and county yields from the expected county yield, which is defined as the average county yield in the previous year. In other words, the yield guarantee for year t is:

$$X_i * \left[\overline{\text{Yield}}_{c,t-1} - \frac{1}{10}\sum_{s=1}^{10} \overline{\text{Yield}}_{c,t-s} + \frac{1}{10}\sum_{s=1}^{10} \text{Yield}_{i,t-s} \right]$$

where X_i is again the coverage level chosen by the farmer.

The final category of insurance plans is group revenue insurance (Group Risk Income Protection or GRIP), which pays farmers based on a combination of county yields and futures prices. As with the individual revenue plans, prices in a county revenue plan are based on CBT futures prices for a given crop. County revenue plan payments are determined as follows:

$$\text{Pay\_GRIP}_{it} = \text{Revenue guarantee}_{ict} - \text{Revenue}_{ct} \tag{9}$$

$$\begin{aligned} \text{Revenue guarantee}_{ict} &= X_i * P\,(pre - season)_t * \frac{1}{T}\sum_{s=2}^{T+1} \widehat{\text{Yield}}_{c,t-s} \\ \text{Revenue}_{ct} &= \overline{\text{Yield}}_{ct} * P\,(harvest)_t \end{aligned}$$

8. Appendix B - Prices and Yields

In this section, I regress the prices of GRP crop insurance on last year's yield, current yield, and expected yields. The level of observation is county-year-coverage level. The coverage level ranges from 65% to 90% in increments of five percentage points.

Table A1 shows the relationship between the price of GRP policies and last year's yield, controlling for expected yield. There is no statistically significant relationship between prices and last year's yield. The expected yield is negatively and significantly correlated with the price.

Table A2 shows the relationship between the price of GRP policies and this year's yield, controlling for expected yield. Again, there is no statistically significant relationship between the current yield and the price. Moreover, the high R-squared suggests that the variables included in the regression capture nearly all the relevant variation in prices.

In Table A3, I show OLS regressions corresponding to Table 3. Although high yields last year are strongly predictive of lower yields this year, there is no significant relationship between last year's yields and the number of currently held policies (although the estimate is not very precise). Current yields are also not significant predictors of the total takeup. For both current and past yields, the sign of the estimated coefficient is the opposite of what would be predicted by theory. Changing the dependent variable to $Log\,(Policies_{lt} + 1)$ to take into account observations in which no GRP policies are purchased does not change the results qualitatively, although it does increase their precision.