8
Willingness-to-pay for ENSO Risk Protection

Even experts struggle to communicate their willingness-to-pay for new or unfamiliar products with many possible configurations, like an ENSO options contract. To manage that difficulty I've augmented chapter 7's qualitative analysis with an adaptive choice-based conjoint analysis - a technique used primarily in the field of marketing to infer individuals' willingness to pay for a product based on their preferences over a small set of product configurations. I obtain estimates of the willingness to pay for various ENSO contracts from a handful of the individual interview subjects. The results of that survey reinforced the message from my interviews (chapter 7): there is demand and supply that will likely cross in new ENSO markets.

To validate the qualitative results from the interview in chapter 7, this brief chapter presents:

- willingness-to-pay for various ENSO contracts from a handful of the interview subjects, estimated via adaptive choice-based conjoint analysis.

This chapter's conjoint analysis is meant to provide quantitative estimates for key parameters that are difficult to approximate through qualitative interviews, such as the elasticities of supply and demand for ENSO risk products. Importantly, respondents revealed their values for those parameters indirectly in the context of choice exercises, similar to those that financial professionals would face should such an ENSO market launch.

Adaptive choice-based conjoint

It is difficult to collect quantitative data on future purchase decisions (e.g. willingness-to-pay (sell) and demand (supply) elasticities). In hypothetical surveys respondents tell researchers that they are willing to pay more for new products than they actually are, as revealed by subsequent purchase decisions. Looking across a wide range of those studies, List and Gallet [2001] and Murphy et al. [2005], offer consensus estimates for the gap between revealed and stated willingness-to-pay ranging from 1.35 to 3 times the underlying real willingness to pay.

Given the poor results from simply asking people what they will pay for a new product or service, many economists and marketers have turned to relative preferences, estimating willingness-to-pay indirectly from the choices that people make in circumstances that are closer to how we actually shop.

Broadly, the set of survey methodologies that back-out preferences from choices is called discrete choice analysis. Discrete choice analysis covers techniques as simple as soliciting yes-or-no reactions to products at a specific price.

It also includes more sophisticated techniques like choice-based conjoint analysis, used for products with a high degree of customizeability. At the heart of conjoint analysis, respondents are given a choice between products that are the result of random attribute sampling. For example a respondent might be asked to indicate their preference between a "red car with cruise control and a sunroof for USD 25,000" and "a black convertible without cruise control for USD 26,000."

After the respondent has answered a series of those questions, it is possible to use logit regression to back out utility measurements for each attribute. By normalizing the coefficients on each attribute (which are given in terms of utility) by the utility of the price coefficient, you can also get willingness-to-pay estimates in dollar terms. Those willingness-to-pay estimates are only valid relative to some benchmark product configuration.

Basic (i.e. *full profile*) conjoint analysis requires large sample sizes because the product configurations are generated entirely by chance. Many survey respondents have to choose among many randomized product configurations to give well-powered estimates of their utility coefficients for each attribute, even when you are estimating those coefficients for a group of people, rather than individuals. That sample size consideration presented me with two problems. First, I did not have a large sample population. There are only a handful of people with the requisite background in derivatives or insurance (preferably with expertise in weather or climate risk), willing to take my survey. Second, I had to be respectful of that population's time. Together, these limitations meant that I simply could not ask the volume of choice questions required by traditional choice-based conjoint analysis.

Instead, with the help of a donation from Sawtooth Software, I used an adaptive choice-based conjoint analysis. Adaptive choice-based conjoint analysis dynamically reconfigures each choice task based on past responses to reduce the redundancy of information revealed by each successive choice. Using those dynamic techniques, Sawtooth offers individual-level utility estimates. Those utilities can be converted into willingness-to-pay/sell estimates and those estimates can be arrayed to form supply and demand curves.

Choice-task methodology

Respondents first identified the level of each attribute that would be part of their ideal risk management tool. That provided a starting point for a multi-dimensional search problem, which the survey design algorithm uses to construct choice questions. Each questions is supposed to narrow the search space for its subsequent questions.

In the second phase of the survey, the software displays attribute combinations that respondents note are either a possibility for them to purchase (or sell, depending on whether they are hedgers or speculators) or not. The algorithm uses these questions to identify attributes that respondents consistently avoid or include in their choices. The algorithm explicitly asks if respondents consider those levels as "Unacceptable" or "Must Have" respectively.

After those first rounds, the algorithm has narrowed the acceptable search space for products of interest to the individual survey respondent. Many comparisons would be redundant for that respondent. For example, choosing between combinations that both contain an "Unacceptable" attribute will not provide much additional information about the relative preferences of that respondent. Attribute configurations that are still in the search space enter a Choice Tournament, where the respondent indicates their top choice among successive displays of three full product configurations. The winning concept from each set of three advances to the next choice set, while its new competitors are assembled from within the remaining search space. The Choice Tournament continues until a winner is determined (a combination that apparently cannot be beat) or a maximum number of questions have elapsed. To the extent that Sawtooth explains its search algorithm, it is available in Johnson and Orme [2007] and Orme [2009b]. Similar algorithms for adaptive conjoint analysis are discussed in greater detail in Toubia et al. [2003] and Toubia et al. [2004].

Estimating utilities through hierarchical Bayesian analysis

Responses to choice tasks are converted into utility scores through a multi-level Bayesian model. Individual utility scores for specific attributes (*parts-worth*) are modeled using a multivariate normal

Feature	Select Feature	Cost for Feature
Underlying index	◯ NOAA NINO 3.4 ◯ NOAA NINO 3 ◯ NOAA NINO 1.2 ◯ Standard weather derivative (ex. heating/cooling degree day) ◯ Multi-peril catastrophe bond index/fund ◯ Standard energy derivative (ex. heating oil or gasoline)	0 BPS
Settlement	◯ Annual ◯ Quarterly ◯ Monthly	0 BPS
Available analysis	◯ Experts (+ 400BPS) ◯ Experts + research reports (+ 350BPS) ◯ Experts + research reports + stress testing tools (+ 300BPS) ◯ Experts + research reports + pricing formulas (+ 50BPS) ◯ Experts + research reports + stress testing tools + pricing formulas	0 BPS
Contract type	◯ Futures (payment based on index value) ◯ Options coving extreme event (payment based on index value in extreme range) ◯ Options coving events of moderate strength or worse (payment based on index value in specific range) ◯ Specialized bilateral transactions (e.g. Insurance-linked securities, swaps, etc.) (+ 100BPS)	0 BPS
Liquidity	◯ Low - ex. most single stock futures (+ 300BPS) ◯ Moderate - ex. CME housing index (+ 100BPS) ◯ Moderate plus - ex. Czech Koruna (+ 50BPS) ◯ High - ex. Platinum	0 BPS
	Total	400 BPS

Figure 8.1: Screen capture of survey - in this task respondents choose the attributes that would be part of their ideal risk management tool.

Here are a few derivatives you might like. For each one, indicate whether it is a possibility or not.

(1 of 8)

Underlying index	Standard weather derivative (ex. heating/cooling degree day)	Standard energy derivative (ex. heating oil or gasoline)	Standard weather derivative (ex. heating/cooling degree day)	NOAA NINO 3
Settlement	Quarterly	Quarterly	Monthly	Annual
Available analysis	Experts	Experts + research reports + stress testing tools + pricing formulas	Experts + research reports	Experts + research reports + stress testing tools + pricing formulas
Contract type	Futures (payment based on index value)	Options coving extreme event (payment based on index value in extreme range)	Options coving events of moderate strength or worse (payment based on index value in specific range)	Options coving events of moderate strength or worse (payment based on index value in specific range)
Liquidity	Moderate plus - ex. Czech Koruna	High - ex. Platinum	Low - ex. most single stock futures	High - ex. Platinum
Expected return in BPS above LIBOR (to the speculator)	597BPS	545BPS	1,572BPS	668BPS
	○ A possibility ○ Won't work for me	○ A possibility ○ Won't work for me	○ A possibility ○ Won't work for me	○ A possibility ○ Won't work for me

UK
UNIVERSITY OF
KENTUCKY

Figure 8.2: Screen capture of survey - in this task respondents indicate whether they consider algorithmically generated product configurations to be a possibility for their use.

Among these three, which is the best option? (I've grayed out any features that are the same, so you can just focus on the differences.)

(1 of 10)

Underlying index	NOAA NINO 3.4	Standard weather derivative (ex. heating/cooling degree day)	Standard energy derivative (ex. heating oil or gasoline)
Settlement	Quarterly	Quarterly	Quarterly
Available analysis	Experts + research reports + stress testing tools + pricing formulas	Experts	Experts + research reports
Contract type	Options coving extreme event (payment based on index value in extreme range)	Options coving events of moderate strength or worse (payment based on index value in specific range)	Futures (payment based on index value)
Liquidity	Moderate plus - ex. Czech Koruna	Moderate - ex. CME housing index	High - ex. Platinum
Expected return in BPS above LIBOR (to the speculator)	330BPS	281BPS	1,226BPS
	○	○	○

UK UNIVERSITY OF KENTUCKY

Figure 8.3: Screen capture of survey - in this choice tournament question respondents identify their preferred risk management tool from among algorithmically generated options.

distribution:

$$\beta_{\text{respondent,attribute}} \sim \mathcal{N}(\mu_{\text{attribute}}, \sigma_{\text{attribute}}) \tag{8.1}$$

Where:

- $\beta_{\text{respondent,attribute}}$ is the utility score that a respondent assigns to an attribute;
- $\mu_{\text{attribute}}$ is the average utility score of all respondents for an attribute, and;
- $\sigma^2_{\text{attribute}}$ is the variance of that utility score (which is part of a matrix, D, of variances and covariances of the distribution of parts-worth across individuals).

Those utility scores are linked to respondents' actual choices through a multinomial logit model. The probability that an individual chooses a certain product given the options in a choice task is denoted $\text{Pr}_{\text{respondent}}(\text{product}_{\text{choice task}})$ and defined as:

$$\text{Pr}_{\text{respondent}}(\text{product}_{\text{choice task}}) = \frac{\exp\left(\sum_{\text{attribute of product}=1}^{\text{max attributes in product}} \beta_{\text{respondent,attribute}}\right)}{\exp\left(\sum_{\text{product in choice task}=1}^{\text{max products in choice task}} \sum_{\text{attribute of product}=1}^{\text{max attributes in product}} \beta_{\text{respondent,attribute}}\right)} \tag{8.2}$$

Parameters β, μ, D are estimated via Gibbs sampling[1].

$\beta_{\text{respondent,attribute}}$ are denominate in utils. However, the choice tasks include price as one variable attribute. So we also get a parameter estimate for $\beta_{\text{respondent,unit price}}$ that gives us a ratio of utiles to price units for a given respondent. In this case, the price units are basis points above LIBOR in expected return.

[1] B Orme. The CBC/HB system for hierarchical Bayes estimation version 5.0 technical paper. Technical report, Sawtooth Software, 2009a. URL http://www.sawtoothsoftware.com/education/techpap.shtml

Normalizing the other parts-worth parameters by that utility of price, gives us a willingness-to-pay/sell denominated in basis points. Note that those willingness-to-pay/sell estimates are only valid relative to some baseline. In this case, I choose a standard multi-peril CAT bond for my baseline. A standard multi-peril CAT bond is indicated in the survey as a "bilaterally negotiated contract, traded at low volume, settling annually on a multi-peril CAT bond index/fund, supported by experts, research reports, and stress testing tools."

Results

Tables 8.1 and 8.2 provide utility scores[2] for hedgers and related ser-

[2] Note that the utility scores are zero-centered across each attribute category.

vice providers ($n = 8$) and speculators ($n = 7$) respectively.

Clearly, this is a small sample. The standard deviations on the utility scores in tables 8.1 and 8.2 reflect that. Very few of the mean estimates are more than two standard deviations away from zero.

Bayesian analysis generally does not use simple thresholds for statistical significance. Instead, parameters with wider posterior distributions are interpreted as subject to greater uncertainty than those with tighter distributions. So, these estimates should not be disregarded, especially given the fact that the sample driving that uncertainty is representative of the small community of people who might be in a position to buy, sell, or advise on ENSO derivatives. However, the evident uncertainty in the estimates does suggest that beliefs formed from these results should be held loosely.

Table 8.1: Average utilities (zero-centered differences) of product characteristics for hedgers and service providers

Characteristics	Average Utilities	StDev
NOAA NINO 3.4	39.70	18.67
NOAA NINO 3	-8.64	51.30
NOAA NINO 1.2	-13.29	52.48
Standard weather derivative (ex. heating/cooling degree day)	8.40	25.09
Multi-peril catastrophe bond index/fund	-11.75	91.00
Standard energy derivative (ex. heating oil or gasoline)	-14.42	29.69
Annual	-22.63	53.66
Quarterly	5.89	18.53
Monthly	16.75	44.30
Experts	-21.70	22.92
Experts + research reports	7.96	23.18
Experts + reports + stress testing tools	3.22	26.78
Experts + reports + pricing formulas	-2.42	17.24
Experts + reports + stress testing + pricing	12.94	29.64
Futures (payment based on index value)	5.45	29.29
Options covering extreme event	26.83	16.41
Options covering events of moderate strength	14.03	33.46
Specialized bilateral transactions	-46.30	62.47
Low - ex. most single stock futures	-29.90	25.68
Moderate - ex. CME housing index	11.41	39.81
Moderate plus - ex. Czech Koruna	-3.04	26.59
High - ex. Platinum	21.54	39.54
E[annual return above LIBOR] = 40 BPS (to the speculator)	51.72	64.93
E[annual return above LIBOR] = 2280 BPS (to the speculator)	-51.72	64.93
None	48.67	27.07

Table 8.1 suggests that hedgers prefer:

- Niño 3.4 to Niño 1.2 or Niño 3;

- monthly contract settlement to annual or quarterly;
- more available analysis for their trades, and are particularly weary of products in which hedging decisions are only supported by expert opinion;
- options on extreme events (they are disinclined toward bilateral transactions);
- more liquid contracts;
- paying less for hedges (i.e. they show higher utility for paying 40 basis points above LIBOR than for paying 2280.)

Table 8.2: Average utilities (zero-centered differences) of product characteristics for speculators

Characteristics	Average Utilities	StDev
NOAA NINO 3.4	23.52	18.41
NOAA NINO 3	-21.40	18.16
NOAA NINO 1.2	-13.59	15.09
Standard weather derivative (ex. heating/cooling degree day)	23.73	22.99
Multi-peril catastrophe bond index/fund	-1.67	53.16
Standard energy derivative (ex. heating oil or gasoline)	-10.58	14.73
Annual	9.63	25.42
Quarterly	2.69	17.43
Monthly	-12.33	25.30
Experts	-33.61	37.08
Experts + research reports	-9.64	25.19
Experts + reports + stress testing tools	5.29	16.06
Experts + reports + pricing formulas	13.63	30.68
Experts + reports + stress testing + pricing	24.33	40.81
Futures	-13.00	34.59
Options covering extreme event	5.15	11.58
Options covering events of moderate strength	2.65	23.60
Specialized bilateral transactions	5.20	26.45
Low - ex. most single stock futures	-12.24	18.60
Moderate - ex. CME housing index	19.38	20.02
Moderate plus - ex. Czech Koruna	-9.32	12.72
High - ex. Platinum	2.18	21.36
E[annual return above LIBOR] = 40 BPS (to the speculator)	-141.69	52.54
E[annual return above LIBOR] = 2280 BPS (to the speculator)	141.69	52.54
None	-63.53	77.47

By contrast, table 8.2 suggests that speculators prefer:

- contracts settling based on standard weather indexes or Niño 3.4 (between which they are largely indifferent) to all other indexes;

- annual to quarterly or monthly settlement, but not by a large margin;
- also to trade with more available analysis;
- options on extreme events and bilateral transactions (between which they are indifferent) to other contract forms;
- moderately liquid contracts;
- to be paid more for hedges.

Most of these inferences are obvious or reinforce what I found qualitatively in my interviews.

Table 8.3: Average importance of product characteristics for hedgers and service providers

Characteristics	Percent	StDev
Underlying index	21.39	12.91
Settlement	12.02	12.40
Available analysis	11.52	5.14
Contract type	17.04	10.98
Liquidity	15.37	3.99
Expected return in BPS above LIBOR (to the speculator)	22.65	14.89

Table 8.4: Average importance of product characteristics for speculators

Characteristics	Percent	StDev
Underlying index	14.85	6.35
Settlement	7.77	4.17
Available analysis	12.33	9.98
Contract type	9.53	4.89
Liquidity	8.29	3.83
Expected return in BPS above LIBOR (to the speculator)	47.23	17.51

In tables 8.3 and 8.4 the variance within utility scores for a given attribute category is compared to the variance across attribute categories to provide a measure of the importance of each product category to hedgers' and speculators' choices[3]. They show that speculators' decisions are driven by price. Hedgers also look first to price, but it is not their only concern. Price is roughly as important to hedgers as other product attributes. Inference about hedgers' price concerns are subject to greater uncertainty than for speculators.

[3] Bryan Orme. Fine-tuning CBC and adaptive CBC questionnaires. Technical report, Sawtooth Software, 2009b. URL http://www.sawtoothsoftware.com/education/techpap.shtml

I converted the raw utility score estimates of individual respondents (not zero-centered as in tables 8.1 and 8.2) into willingness to pay/sell relative to a standard CAT bond. By taking the empirical cumulative distribution function of those individual estimates of any given product configuration (and taking the inverse of that function for hedgers), I put together supply and demand curves for various product designs.

The supply curve, for example, shows what percentage of the surveyed speculators would be willing to sell at an expected return denominated in basis points over LIBOR, relative to a standard multi-peril CAT bond.

Figure 8.4 displays those curves for monthly settled Niño 3.4 contracts. Each column of graphs shows the curves for a different contract type. Each row corresponds to a different level of supporting analysis available to hedgers and speculators. Within each graph, colors distinguish supply and demand curves for hedgers and speculators operating under different liquidity levels. For example, the green line in each graph shows a demand curve for monthly settled Niño 3.4 protection with high liquidity.

One clear lesson from figure 8.4 is that speculators have remarkably uniform willingness-to-sell. Almost regardless of the underlying product configuration, they will sell Niño 3.4 protection at rates ranging from a bit above a standard CAT bond to a bit below. As we saw in table 8.4 the elasticity of supply is high, with risk managers responding to expected returns above all else.

Froot [1999] discusses monopoly pricing power in catastrophic risk markets. Given that possibility, figure 8.4 provides some cause for concern. The most motivated hedgers in the survey appear price insensitive. This suggests that if one or two risk sellers were able to isolate the most motivated hedgers, and avoid creating a competitive market for ENSO protection, they may enjoy some pricing power.

Fortunately, figure 8.4 also shows that most hedgers are sensitive to price and that the market will clear in that elastic region of the demand curve. At the clearing price (just above the CAT bond average), most of the hedgers in the sample would have coverage.

That accords roughly with my interviews, which suggest that prices in the ENSO market will be linked to those in the CAT bond market. Unlike the survey, however, my interviews suggested that ENSO will be on the low end of the CAT bond pricing spectrum. The survey also indicates that there will be risk taking capacity to meet virtually all demand for ENSO coverage. It is encouraging to see that respondents' opinions do not change fundamentally after being filtered through a sophisticated preference elicitation routine.

Conclusions

This final chapter supports the qualitative findings in chapter 7. It presents the results from an adaptive choice-based conjoint analysis, arranged to form indicative supply and demand curves for a host of ENSO risk management products. Those supply and demand curves suggest that there is latent demand and supply that will cross. The

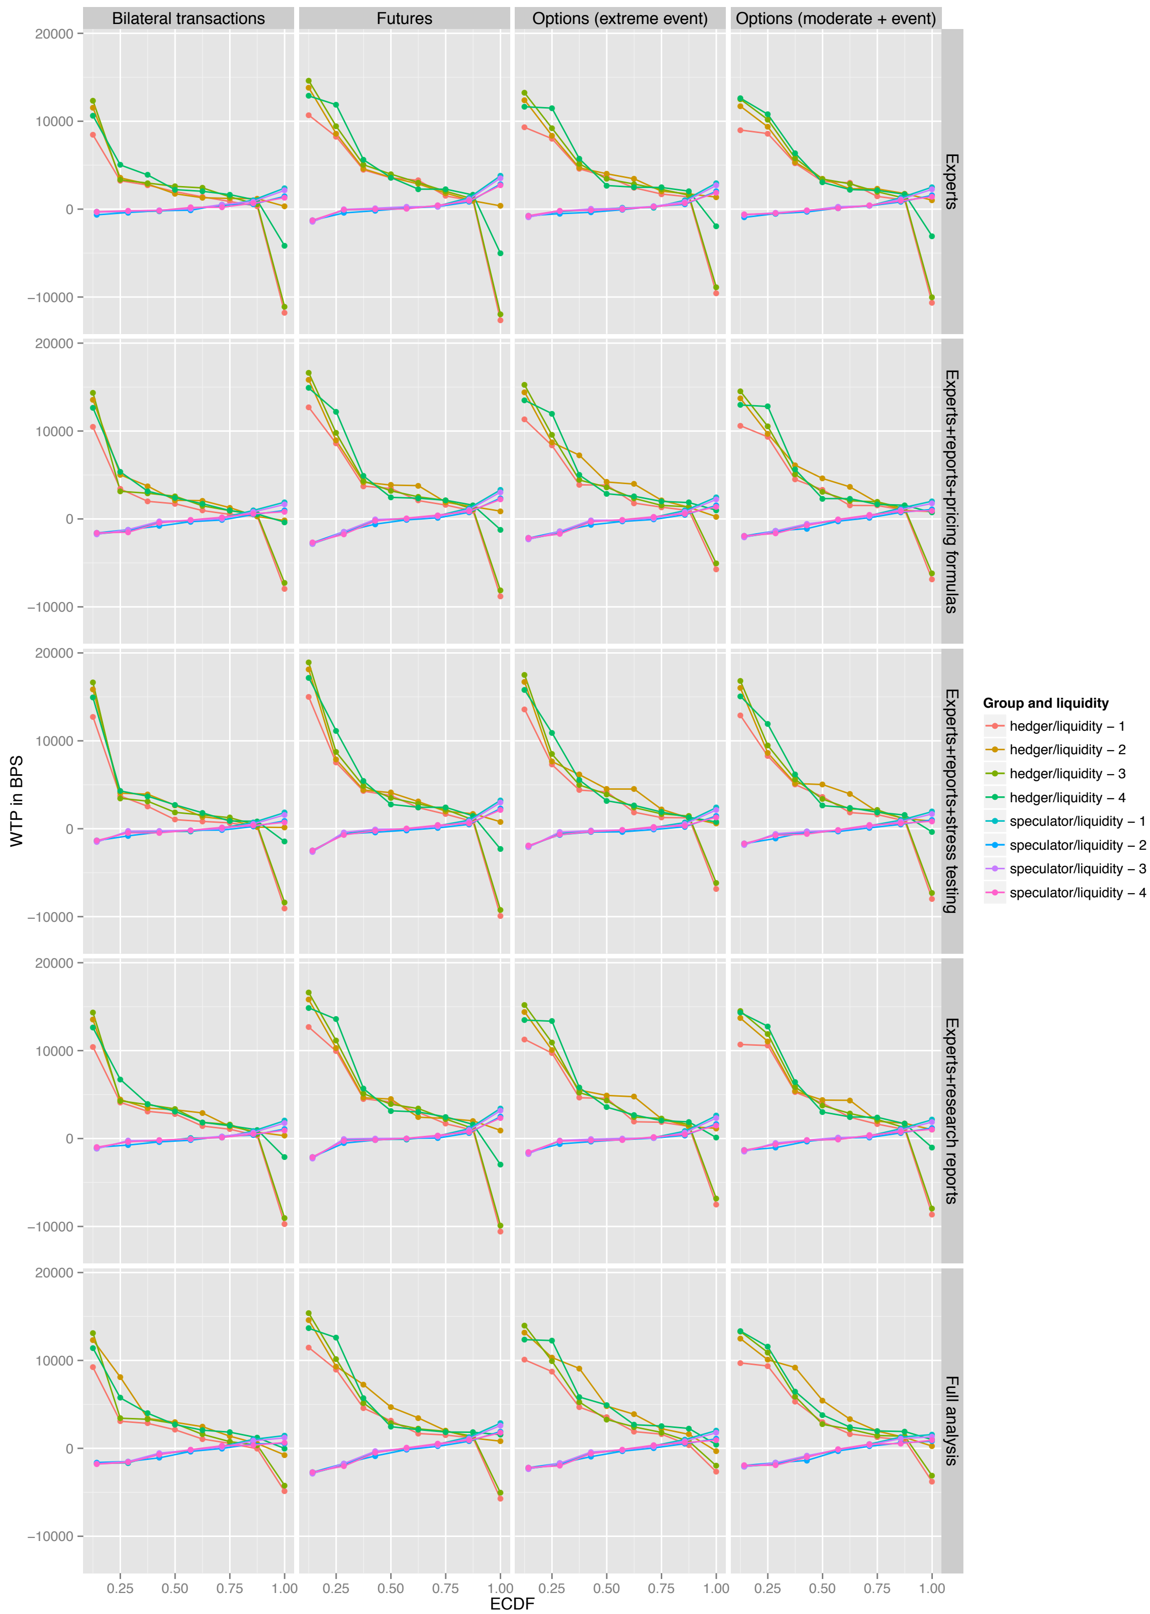

Figure 8.4: Relative supply and demand curves for monthly settled Niño 3.4 contracts. Willingness-to-pay/sell is in basis point relative to a standard multi-peril catastrophe bond. See figure 6.7 for recent benchmark prices for catastrophe bond risk. Estimates are based on the raw utility scores of individuals in an adaptive choice-based conjoint survey. The

market will clear at prices that offer speculators expected returns similar to those from multi-peril CAT bonds. That is above the markups indicated by speculators in the interviews profiled in chapter 7. Additionally, the curves suggest that there will be elastic demand at the market clearing price. That means that speculators will have limited opportunity to exercise pricing power.

In general, speculators were willing to offer any risk management contract, given a high expected return on the underlying risk. By contrast, hedgers' purchase decisions will take into account a wider array of considerations, including the contract type (ILS, futures, options, etc.). As expected, hedgers have stronger preferences for specific underlying indexes than speculators.

These findings are subject to substantial uncertainty, given the size of the underlying sample. However, they provide important validation of the interviews in chapter 7. Past studies suggest that qualitative interviews have been unreliable in revealing willingness-to-pay for unfamiliar products types. Researchers have documented better predictive insight from the type of choice-based survey used in this chapter.

Together, the chapters suggest that ENSO markets will be supported by latent demand and supply sufficient to support formal risk management trades in the near-term.