

## A Computational Theory of Willingness to Exchange

Pete Lunn and Mary Lunn<sup>a</sup>

*Abstract:* A new model of exchange is presented following Marr's conception of a "computational theory". The model combines assumptions from perceptual theory and economic theory to develop a highly generalised formal model. The approach departs from previous models by focussing not on how ownership alters preferences, but instead on difficulties inherent in the process of exchange in real markets. Agents treat their own perceptual uncertainty when valuing a potential exchange item as a signal regarding the variability of potential bids and offers. The analysis shows how optimising agents, with no aversion to risk or loss, will produce an endowment effect of variable degree, in line with empirical findings. The model implies that the endowment effect is not a laboratory finding that may not occur in real markets, but rather a market phenomenon that may not occur in the laboratory.

*Keywords:* Endowment effect, willingness to accept, willingness to pay, exchange, uncertainty

*Corresponding Author:* [Pete.Lunn@esri.ie](mailto:Pete.Lunn@esri.ie)

---

This work is supported by a Research Development Initiative grant to the first author from the Irish Research Council for the Humanities and Social Sciences (IRCHSS). The authors are grateful to Daniel Goldstein, David Duffy, Denis Conniffe, Liam Delaney, Laura Malaguzzi Valeri and Iván Barreda-Tarrazona for helpful comments on earlier drafts of this paper. We are also grateful to various attendees at the Royal Economic Society conference 2010, the 2<sup>nd</sup> Economic Psychology Ireland conference and seminar participants at the Economic and Social Research Institute and NUI Galway, for helpful input.

<sup>a</sup> Department of Statistics, University of Oxford, OX1 3TG, UK.

ESRI working papers represent un-refereed work-in-progress by researchers who are solely responsible for the content and any views expressed therein. Any comments on these papers will be welcome and should be sent to the author(s) by email. Papers may be downloaded for personal use only.

# A Computational Theory of Willingness to Exchange

## 1. Introduction

Exchange of ownership is a fundamental economic process, yet the willingness of people to make exchanges remains, at best, partially understood. This paper offers an alternative and novel theoretical treatment of willingness to exchange, which aims to combine the methodological approach to psychology and neuroscience of David Marr (1982) with the traditional optimisation framework of economics. This blend of psychological and economic approaches results in a formal and highly generalisable model of willingness to exchange, which may offer insight into at least some of the empirical literature on exchange of ownership and, in particular, the classic “endowment effect” experiments.<sup>1</sup>

Marr (1982) argued that an information processing system could be understood at three distinct analytical levels: computational, algorithmic and physical. This tri-level approach became rapidly influential in the study of perception and, ultimately, throughout much of cognitive psychology. Our focus is on the computational level, at which researchers seek “functional descriptions of what information processing systems, including brains, are designed to do” (Marr, 1982). Computational theories explain not only what the system does, but why it makes sense for it so to do. This distinction has similarities with those between normative and positive economics and between adaptive and non-adaptive traits in evolutionary science. We prefer to consider the theory of willingness to exchange we present as a computational theory in Marr’s sense, however, since it is not intended to be normative and does not explicitly model an evolutionary process. Nevertheless, the theory does more than describe observed behaviour; it proposes a rationale for it.

Adopting Marr’s approach, we argue that the computational problem faced by economic agents seeking to exchange items (including monetary amounts) consists of three essential aspects, the first two of which are well supported by bodies of empirical evidence (see Subsections 2.2 and 2.3). First, agents must assess the private values of the items in question, which may be subject to significant uncertainty. Second, because price dispersion is a ubiquitous feature of product markets, agents must take account of variability in the distributions of bids and offers. Third, agents must weigh up potential surpluses against the accumulating cost of the time and/or effort required to obtain, consider and act on bids and offers. The computational problem for agents aiming to generate surplus through exchange is, therefore, simultaneously to weigh up the perceived value of the good, the perceived distribution of likely bids and offers, and the accumulating cost of attempting exchange.

---

<sup>1</sup> Although in widespread use, the expression “endowment effect” has become controversial, as some researchers argue that it assumes the cause of the empirical phenomenon in question (Plott and Zeiler, 2005). Here, it is used only to refer to the empirical phenomenon of unwillingness to make apparently beneficial exchanges, manifested either through a disparity between buying and selling prices (Kahneman, Knetsch and Thaler, 1990) or reluctance to exchange an endowed good for another that is not owned (Knetsch, 1989).

From this framework, we derive a formal model of optimal exchange and several consequent propositions, including one that offers a potential heuristic for learning to set optimal willingness to accept (WTA) and willingness to pay (WTP) in real markets. The central argument made is that the endowment effect may at least in part result from the extension of this behaviour, which is well adapted for real markets, into experimental markets with only one opportunity to trade at one price, where the behaviour appears to be suboptimal. The theory is novel and contrasts with pre-existing models applied to the empirical phenomena of ownership and exchange. Specifically, it does not centre on the shape of preference functions, but on challenges inherent in the process of exchange. As a result, our model predicts a substantial WTA-WTP gap between without making any assumptions about preferences over risk or loss. It is consistent with some empirical findings and also generates new testable hypotheses.

Our approach is unashamedly deductive and theoretical and is thus most similar to the recent contributions of Loomes, Orr and Sugden (2009) and Isoni (2011). The aim is to provide an alternative theory of exchange for consideration alongside prevailing theories, which presently struggle to account for the now very substantial spectrum of empirical findings (Ericson and Fuster, 2013). We make no claim that our computational model of willingness to exchange can explain all of these relevant phenomena, but it is consistent with many of the main replicable findings associated with the endowment effect. Moreover, we have successfully tested novel predictions of the model, which are not consistent with models based on loss aversion (Lunn and Lunn, 2014).

Section 2 supports and motivates our assumptions, drawing on findings from economics and psychophysics. Section 3 presents the model. Section 4 relates the model to empirical findings on exchange. Section 5 concludes.

## **2. Model Assumptions**

As outlined above, our model begins by considering exchange to be a computational problem of information processing. Agents must assess the private values of items to be exchanged, compare these with a perception of the distribution of potential bids and offers in the market, and weigh these up against the accumulating costs of ongoing interactions with more potential trading partners. Framed in this way, the problem is to convert the available perceptual information into a decision-rule about when to trade, in order to secure surplus. The present section expands on this logic, employing empirical evidence where necessary, to substantiate the assumptions underlying the formal model.

### *2.1 Relationship to Other Models*

Space does not permit a thorough review of the now large number of theories and models designed to explain aspects of willingness to exchange, so here we briefly explain how our approach contrasts it with extant theories.

The primary approach taken by researchers seeking to explain the endowment effect and related findings relating to willingness to exchange is to assume that economic agents are essentially loss averse (Tversky and Kahneman, 1991; Köszegi and Rabin, 2006; Loomes et

al., 2009). The primary challenge facing such accounts is to explain why the extent of loss aversion varies so much with experimental procedures (e.g. Franciosi et al., 1996; Shogren et al, 1994; Plott and Zeiler, 2005, 2007), trading experience (e.g. List, 2003, 2004) and the type of item being exchanged (Horowitz and McConnell, 2002; Sayman and Öncüler, 2005; Bateman et al., 2005). The present approach is distinct, because it does not centre on the shape of preferences over goods and money. Instead, our computational theory of willingness to exchange starts with the idea that agents seeking to exchange must consider not only what items are worth to them, but also what they are worth to others.

Two recent models do incorporate this possibility (Isoni, 2011; Weaver and Frederick, 2012). In both cases, the assumption is that agents get utility both from consuming the item and from the transaction itself, such that good deals result in higher overall utility. Our approach again contrasts, because we foresee no role for transaction utility. Instead, the distribution of valuations of others determines optimal willingness to exchange because it dictates the probability of making a trade, which agents in real markets must routinely take into account.

Because the model presented here does not involve a change in preferences resulting from ownership, it is also distinct from various theories that have invoked particular psychological mechanisms of preference formation, most of which have focussed on why owning an item might increase its private value (e.g. Beggan, 1992; Johnson, Haubl and Keinan, 2005; Ashby, Dickert and Glöckner, 2012; Bordalo, Gennaioli and Shleifer, 2012). It is important to note that our theory does not rule out the possibility that preferences are changed by ownership. Rather, we raise the possibility that some of the main findings relating to willingness to exchange might be governed by other influences. Lastly, our model has some similarities with that of Kling, List and Zhao (2010), who also propose a dynamic mechanism to explain willingness to trade, by placing agents in a context where there is potentially more than one opportunity to trade. Their model still assumes that ownership alters a preference, such that sellers place a higher option value on future selling through cognitive dissonance. Again, our model does not require such a mechanism.

## *2.2 Exchange Involves Perceptual Discrimination*

Our approach instead begins with perceptual theory. When deciding whether to trade one item for another an agent must, at a minimum, discriminate which item has the greater private value. In other words, the agent must perceive the value of what is potentially obtained and compare it to a perception of the value of what is given up, or similarly compare perceptual representations from memory. An exchange, therefore, has parallels with a two-alternative forced-choice discrimination task of the sort routinely used by psychophysicists to measure the precision of human perception. The analogy may be instructive. While orthodox micro-economic theory assumes stable and well-ordered preferences, psychophysical evidence suggests that private valuation is likely to be a complex and approximate perceptual task. Valuing a consumer good is likely to involve immediate perceptions of multiple attributes of the item, perceptions of past experiences of the item held in memory, and perhaps additional indirect sources information about the item's quality. Yet, ultimately, agents must combine these inputs into an internal representation of value to compare against different items or against monetary amounts.

Given this, the existing literature in perceptual science raises two questions. First, to what extent are internal representations of private value subject to error? Second, if private values are indeed approximate, are people able to take the extent of error into account in judgements?

There is relevant psychophysical evidence on both counts. The precision of perceptual discrimination is often measured via the Weber fraction – the proportionate increase/decrease in a stimulus that is required before it can be reliably perceived as larger, louder, heavier etc. The smallest Weber fractions obtained for the discrimination of basic perceptual dimensions, e.g. size or shape, are around 2 – 8% (e.g. Burbeck, 1987), i.e. a stimulus must be 2 – 8% greater on the given dimension to be reliably perceived as such. Weber fractions for the weight of an object held in the hand are typically around 10% (e.g. Brodie and Ross, 1984). For more complex properties of three-dimensional shapes, Weber fractions may be as high as 15 – 30% (Lunn and Morgan, 1997). Thus, perceptions of even basic visual and haptic dimensions are subject to significant error that varies systematically by task. Given these findings, the extent of error when valuing consumer items is likely to be considerable. Valuation requires an agent to judge not only perceptual primitives like size and weight, but also properties like attractiveness and durability, before even taking account of perceptions of socially influenced factors that affect value, such as fashionability. How accurately people can perceive such properties and combine them into an overall estimate of private value is not known. However, given the magnitude of error surrounding perceptual primitives, our internal representations of value are likely to be subject to significant error, reflecting fundamental limitations of human perception that are common to us all.

Some previous studies have raised the issue of uncertainty of value in relation to exchange. Loomes et al. (2009) in particular incorporate “taste uncertainty” into their model, such that agents take into account multiple future “taste states”. Our concept of perceptual error is distinct, because we see the uncertainty not as due to variable tastes, experiences or moods, but as a fundamental limitation on the precision of human information processing. This distinction is crucial, because it implies that the extent of perceptual error when valuing different goods will be correlated across individuals. While different people may have different expertise in assessing the private value of exchange items, there will also be commonality. Some items are simply harder to value than others, whatever one’s tastes.

Turning to the second question, there is indeed evidence that people can estimate the extent of their own perceptual error and incorporate this source of uncertainty into judgements. For instance, it is possible to compare performance in assessing the shape of an object via vision, via touch and simultaneously via both vision and touch, where the observer must combine information from both senses to reach a judgement. In performing such tasks, people weight the information from the two different perceptual systems inversely according to the degree of error associated with each perceptual system, as measured by performance in the tasks using vision or touch alone. Subjects’ judgements resemble the outcome of maximum likelihood estimation (Ernst and Banks, 2002).

Given the above, we build three psychophysical assumptions into our computational theory of willingness to exchange. First, perceptual error when valuing goods is likely to be considerable. Second, the extent of error varies systematically with the type of good, such that it is correlated across individuals. Third, people are able to take account of this type of uncertainty in their decisions.

### *2.3 Price Dispersion*

Real markets involve considerable price dispersion and human exchange behaviour is likely to have adapted to cope. With reference to modern markets, Baye, Morgan and Scholten (2006) review more than 30 empirical studies spanning a century of data and a large variety of consumer products. While there is great variation in the extent of dispersion, it is generally substantial and price differences in excess of 30% are common even in competitive markets for homogeneous consumer goods, with no indication that price dispersion has diminished in modern times, despite near costless price comparison in some markets.

Since price dispersion is and remains a characteristic of real markets, agents looking to exchange must routinely perceive and take account of the distribution of potential bids and offers in each market. This is, again, a difficult information processing problem. We are not aware of studies that have directly measured the accuracy of perceptions of the price distribution in real markets. However, studies of consumers' abilities to judge market prices for products (Matthews and Stewart, 2009) or the willingness to exchange of others (Frederick, 2012) imply considerable inaccuracy, both in terms of systematic bias and variability. Given this mix of findings from economics, consumer research and psychology, we assume that agents face price dispersion but differ in their estimates of the distribution of potential bids and offers.


To the three psychophysical assumptions and the assumption of considerable variation in potential bids and offers, we add the uncontroversial assumption that trading activity itself is costly. It takes time and effort to consider and execute exchanges.

## **3. Model**

### *3.1 Optimal Willingness to Accept (WTA)*

The basic set-up is depicted in Figure 1, where we consider an agent who must decide the minimum price they will accept for an item owned (WTA), in order to determine rejection or acceptance of bids. The model can be adapted easily to one in which the agent posts a selling price, or for the pure exchange of items rather than items for money, but we confine ourselves to deriving WTA and WTP for simplicity and because empirical studies of the endowment effect routinely elicit values for WTA and/or WTP.

Figure 1: Model set-up for willingness to accept (WTA). The agent owns good of perceived value  $X$ , perceives sequential future bids of value  $Y$ , each costing  $c$ , and aims to set WTA to maximise surplus.



We assume that an agent's internal perceptual representation of the value of each good consists of a continuous probability density function over a range of possible values. Perceptual error is considerable, such that variability is relatively large with respect to expected value. The agent is endowed with a good, which they perceive to be of value  $X$  with mean  $\mu_x$  and variance  $\sigma_x^2$ . The agent perceives the distribution of bids they can obtain for the good. We make the simplifying assumption throughout that bids correspond to quantities of money, the values of which are represented without error. Moreover, we assume that the mapping of perceived value to numerical amounts is perfect, such that WTA itself is represented without error. Neither simplifying assumption will hold in reality, but we anticipate that the variabilities involved are small relative to the perceptual error that underpins the model. Thus, we can directly compare the perceived value of the good,  $X$ , with the perceived value of future bids,  $Y$ , with mean  $\mu_y$  and variance  $\sigma_y^2$  on the same dimension of value. In Figure 1, the distributions of  $X$  and  $Y$  are normal, while  $\mu_y > \mu_x$  and  $\sigma_y > \sigma_x$ , but these properties are unnecessary for our main result. Based on these perceptions, the agent sets their minimum acceptable price for selling the good,  $WTA = \mu_x + \alpha$ . We assume that the agent aims to maximise expected surplus.

Based on these perceptions, the agent sets their minimum acceptable price for selling the good,  $WTA = \mu_x + \alpha$ . We assume that the agent aims to maximise expected surplus.

Given commonalities of perceptual limitations described in Subsection 2.2, the extent of perceptual error surrounding the agent's valuation,  $\sigma_x^2$ , is positively correlated with the perceptual error of other market participants and the agent takes this correlation into account in their representation of the distribution of potential bids. The variability in the agent's perception of likely bids is also determined by other factors, such as their perception of variation in tastes or needs, such that  $\sigma_y^2 = f(\sigma_x^2, \tau_x^2)$  where  $\tau_x^2$  captures the perceived variability that is not due to perceptual error, with  $f_{\sigma_x^2} > 0$  and  $f_{\tau_x^2} > 0$ . The relationship between  $\mu_x$  and  $\mu_y$  depends on whether the perceived value of the good to the agent is more or less than the value they perceive it to have for others, and on the degree of surplus they expect bidders to build in to bids. It is not necessary to constrain either to obtain our results; nor is it desirable, since we are seeking a highly generalised result.

We assume that the agent expects to receive bids in sequence  $\{Y_1, Y_2, \dots, Y_i, \dots\}$  and sets WTA in advance of receiving bids. Clearly, it is possible for the agent to update their perception of the distribution of bids in light of the ongoing sequence of bids they receive, but for the present we do not incorporate updating in the model. We assign a (small) cost,  $c$ , to receiving each bid, which we call the "encounter cost". One way to conceive of the encounter cost is that the number of encounters in which bids are received determines the length of (costly) time it takes to make the sale, although other conceptions are possible, including equating it to a search cost. We assume, for the present model, that the encounter cost is perceived accurately.

Given these assumptions, we can formulate the expected surplus. Assuming a sale is made, the agent expects to receive a price equal to the expected bid given that the bid is greater than WTA,

$$E(\text{price}) = E(Y | Y > \mu_x + \alpha) \quad (1)$$

In addition to giving up the good, the agent incurs encounter costs as a result of rejecting bids. In setting WTA, they determine a fixed probability of accepting a bid. Thus, the expected number of encounters required to make a sale conforms to a geometric distribution, with parameter  $\Pr(Y > \mu_x + \alpha)$ . The expected total encounter cost (up to and including making the sale) is therefore given by the encounter cost multiplied by the reciprocal of the probability of making a sale,

$$E(\text{total encounter cost}) = \frac{c}{\Pr(Y > \mu_x + \alpha)} \quad (2)$$



Thus, the agent can be considered to face an optimisation problem, in which the aim is to maximise the expected surplus,  $E(S)$ , from selling the good. Combining equations (1) and (2), the agent chooses  $\alpha$  to maximise

$$E(S) = E(Y | Y > \mu_x + \alpha) - \mu_x - \frac{c}{\Pr(Y > \mu_x + \alpha)} \quad (3)$$

This optimisation problem is the heart of the model and is a computational theory in exactly Marr's sense: a theory of what the information processing system does and why it makes sense so to do. Equation (3) represents a trade-off. Increasing  $\alpha$  increases the expected price, but also increases expected encounter costs.

This model structure shares features with some models of consumer search, perhaps most notably that of Reinganum (1979), but while the similarity is instructive there are major differences too. Our model of exchange is much more general. It applies to selling as well as buying. The encounter cost need not be a search cost. Moreover, while consumer search models are concerned with deriving an equilibrium between the consumer's optimum search strategy and the firm's optimum pricing strategy, we are concerned with deriving adaptive buying and selling strategies for exchange, involving generalised distributions of perceived value, bids and offers, whether there are firms involved or not. Contrary to most search models, we do not assume homogeneous buyers (or sellers) with perfect information regarding the distribution of offers (bids). Indeed, we consider this assumption unrealistic and instead conjecture that forming a perception of the distribution of offers (bids) is crucial to determining WTA (WTP). With respect to the market, all that is necessary for our results is that there is some price dispersion linked to the extent of perceptual error in valuation.

The solution to the optimisation problem is derived for any continuous distribution in the Appendix. The existence of a positive  $\alpha^*$  that satisfies (3) depends on

$$c < \int_{\mu_x}^{\infty} (1 - F(y)) dy \quad (4)$$

where  $F(y)$  is the cumulative distribution function of  $Y$ . The condition specified by (4) makes sense: if the encounter cost is too high then there is no price at which a surplus is likely to be made. Assuming the condition in (4) is met, WTA is determined by  $\alpha^*$  which satisfies

$$c = \int_{\mu_x + \alpha^*}^{\infty} (1 - F(y)) dy \quad (5)$$

**PROPOSITION 1:** *For good  $X$  of expected value  $\mu_x$  subject to sequential bids,  $Y$ , each received at a small encounter cost  $c$ , with continuous cumulative distribution function  $F(y)$ , there exists  $\alpha^*$  such that willingness to accept,  $\mu_x + \alpha^*$ , maximises the expected surplus from exchange.*

Given (5),  $\alpha^*$  is decreasing in  $c$ . The higher the encounter cost the greater the need to obtain a sale from fewer encounters. The relationship between  $\alpha^*$  and  $\sigma_y$  is less straightforward,

because it depends on the shape of the distribution of  $Y$ . In the Appendix we derive the following:

PROPOSITION 2: *For bids,  $Y'$ , with consistently higher probability  $Pr(Y' > y) > Pr(Y > y)$  for all  $y$  greater than a fixed value,  $k$ , an agent who maximises surplus will set a higher willingness to accept, such that  $\alpha^{*'} > \alpha^*$ .*

Expressed more simply, the fatter the upper tail of the perceived distribution of bids, the greater  $\alpha^*$  and hence the higher WTA. In all practically applicable cases, where the perceived distribution of bids is unimodal and continuous, with a steadily decreasing probability of receiving bids ever higher than the mean,  $\alpha^*$  is increasing in  $\sigma_y$ . That is, WTA is higher the greater the perceived variability of potential bids.

We further derive the maximum expected surplus from setting WTA according to (5) so as to maximise (3), which the solution:

$$\max E(S) = \alpha^* = WTA - \mu_x \quad (6)$$

Interestingly, for (6) to hold, it must also be the case that

$$E(\text{total encounter cost}) = E(\text{price}) - (\mu_x + \alpha^*) \quad (7)$$

PROPOSITION 3: *The expected total encounter costs of an agent who sets willingness to accept (WTA) to maximise surplus are equal to the expected price over and above WTA.*

From a mental accounting perspective, this solution regarding the expected surplus for optimal WTA is notable. Most straightforwardly, it gives a ready indication of the expected surplus, which may be of benefit to an agent involved in repeated trading activity. Perhaps more importantly, the relationship specified in (7) between the expected price and the expected cost of bids may have heuristic value in helping to set WTA through experience. Equation (3) represents a complex optimisation problem and we make no claim that human subjects internally represent and solve equation (5) for  $\alpha^*$ . What matters here is Marr's (1982) distinction between the computational and algorithmic level of analysis. Equation (5) is a description of the optimal solution to the problem the agent faces, what the agent might aim to achieve, but many potential algorithms might be used to approximate this optimal solution, including relatively simple heuristic ones. Through repeated buying and selling, equation (7) shows that agents will get feedback that is suggestive of setting WTA too high or too low. If a seller repeatedly incurs higher encounter costs than the additional price they ultimately obtain, over and above WTA, then WTA is being set too high, and vice-versa. This is intuitively appealing. We are all familiar with the feeling that it took so long (or required so much effort) to obtain a deal that we regret holding out for the better price, yet there are also times when a quick sale close to the minimum acceptable price leads us to wonder whether we shouldn't have held out for more. Equation (7) shows that if agents balance the time and effort against the additional price obtained above WTA, they are optimising surplus.

Having solved (3) for the general case of a continuous distribution, we now consider the application to the normal distribution, which offers greater insight into the properties of the solution. For  $Y \sim N(\mu_y, \sigma_y^2)$ , the solution to the optimisation problem (see Appendix) is such that


$$\frac{c}{\sigma_y} = \frac{1}{\sqrt{2\pi}} e^{-\frac{\lambda^2}{2}} - \lambda(1 - \Phi(\lambda)) \quad (8)$$

where  $\Phi$  is the cumulative distribution function of the standard normal distribution and

$$\lambda = \frac{\mu_x - \mu_y + \alpha^*}{\sigma_y} \quad (9)$$

Figure 2 shows the relationship between  $\lambda$  and  $c/\sigma_y$ , which is not intuitively obvious from (8) and (9).  $\lambda$  is increasing in  $\sigma_y$  and decreasing in  $c$ . Since  $\sigma_y^2 = f(\sigma_x^2, \tau_x^2)$  and  $f_{\sigma_x^2} > 0$ , the greater the degree of uncertainty in the perception of value, the greater WTA, while the higher the cost of encounters, the lower WTA.


**Figure 2: Optimal willingness to accept (WTA) when perceptual error in valuation is normally distributed. WTA (which is increasing in  $\lambda$ ), is a decreasing function of the cost of sequential encounters in the marketplace,  $c$ , and an increasing function of the perceived variability of future bids,  $\sigma_y$ .**



### 3.2 Optimal Willingness to Pay (WTP)

We now consider the case of WTP, which is depicted in Figure 3. This time, the agent must set WTP for the good of perceived value  $X$  with mean  $\mu_x$  and variance  $\sigma_x^2$ . We assume a perceived distribution of offers with mean  $\mu_z$  and variance  $\sigma_z^2$  and that the agent receives a sequence of offers  $\{Z_1, Z_2, \dots, Z_i, \dots\}$ , each at an encounter cost,  $c$ .

**Figure 3: Model set-up for willingness to pay (WTP).** The agent considers the purchase of a good of perceived value  $X$ , perceives sequential future offers of value  $Z$ , each costing  $c$ , and aims to set WTP to maximise surplus.



The agent's optimisation problem is to choose  $\beta$  to maximise

$$E(S) = \mu_x - E(Z | Z < \mu_x - \beta) - \frac{c}{\Pr(Z < \mu_x - \beta)} \quad (10)$$

Similarly to the solution in the case of WTA,  $\beta^*$  satisfies (see Appendix)

$$c = \int_{-\infty}^{\mu_x - \beta^*} F(z) dz \quad (11)$$

and the expected surplus, given  $\beta^*$ , is given by

$$E(S) = \beta^* = \mu_x - WTP \quad (12)$$

Thus, the general solution is symmetrical to that for WTA and the equivalent of propositions 1 – 3 hold for WTP also. Thus, the greater the degree of perceptual error in valuation, the higher  $\beta$  and the lower WTP, while the higher the cost of encounters, the lower  $\beta$  and the higher WTP.

For the normal distribution we obtain:

$$\frac{c}{\sigma_z} = \eta \Phi(\eta) + \frac{e^{-\frac{\eta^2}{2}}}{\sqrt{2\pi}} \quad (13)$$

where

$$\eta = \frac{\mu_x - \mu_z - \beta^*}{\sigma_z} \quad (14)$$

This solution is again symmetrical to that for WTA, such that  $\eta$  is decreasing in  $\sigma_z$  and increasing in  $c$ .

### 3.3 Model Summary

Before relating the model to the empirical phenomena of exchange, it is important to note how general our solutions are and how few constraints are involved. Given the starting assumptions, the result that optimal WTA (WTP) is increasing (decreasing) with uncertainty in perceived value requires only that the encounter cost is small relative to the variability in valuations. If so, a gap between WTA and WTP will be a characteristic of optimising traders in a market with price dispersion, whether preferences are changed by ownership or otherwise.

## 4. Relationship to Exchange Phenomena

### 4.1 The Endowment Effect

At this stage, the volume of empirical evidence surrounding the endowment effect and closely related phenomena is vast. No one model or theory can hope to account uncontroversially for the stock of findings and it is clearly possible that the effect is the result of multiple influences (see Ericson and Fuster, 2013, for a recent review). This section therefore concentrates on relating the model presented to those findings that have proven robust to replication or that have been clearly identified by systematic reviews or meta-analyses.

Our theory implies a specific interpretation of the standard laboratory result (Knetsch, 1989; Kahneman et al., 1990). If subjects respond to the experimental environment as if they are engaging in typical exchange outside the laboratory, they will instinctively set WTA higher

than WTP or, equivalently, refuse exchanges of owned items for non-owned items they might prefer in a binary choice, because a larger surplus might be possible from an alternative later exchange. The implication is that the endowment effect appears irrational only because the experimental market is a one-shot game with no price dispersion. Real markets are not one-shot games with final outcomes that depend on whether a single announced price exceeds a threshold, or where immediate failure to trade entails permanently lost opportunity. A behaviour that has successfully adapted to real markets may appear irrational in an artificial one-shot/one-price laboratory market that terminates instantly at an exact price.

Some of the experimental manipulations that appear to reduce or remove the endowment effect are consistent with this account. Manipulations that emphasise the one-shot nature of the experiment, or that induce subjects to treat their situation as a single choice rather than a routine opportunity to trade, are inclined to attenuate the effect. Perhaps the most straightforward illustration of this is the substantial variation in the WTA-WTP gap across different value elicitation procedures and experimental instructions (Horowitz and McConnell, 2002; Sayman and Öncüler, 2005). To take a simple example, the WTA-WTP gap reduces when references to “buying”, “selling” and “price” are simply replaced with references to “choosing” (Franciosi et al., 1996). A similar process might help to explain some of the controversial findings of Plott and Zeiler (2005, 2007), partly replicated by Isoni, Loomes and Sugden (2011). Almost all of the numerous experimental manipulations involved in these studies, especially training in the logic of the BDM value elicitation mechanism (Becker, DeGroot and Marschak, 1964), are likely to break the link between behaviour in the laboratory and behaviour in real markets. In the monetary exchange experiments (Plott and Zeiler, 2005), experimenters pointed out to subjects, at length and with examples, that a WTA-WTP gap would lead to inferior outcomes in the experimental market (the standard one-shot market without price dispersion). Similarly, in the non-monetary exchange experiments (Plott and Zeiler, 2007), the “full set of controls” condition that removed the endowment effect replaced the usual practice of inviting subjects to trade the good they owned for another good, and employed instead a decision form asking them to circle the item they wished to take home. A theory of exchange in real markets will apply only if experimental procedures lead subjects to behave as they would outside the laboratory, either because that is their instinctive response to the procedures, or because that is what they believe they are being asked to do. Our interpretation of these controversial findings is straightforward: if subjects believe they are being invited to make a once-off choice, they will choose what they think they prefer, but if they behave as if in a normal market, they will be inclined to set WTA well above WTP and to decline apparently beneficial once-off exchanges of items. Note that this is very different from Plott and Zeiler’s own interpretation that “subject misconceptions” lie behind the endowment effect, since it implies that the original experimental design, in which the endowment effect is large, provides a better indication of willingness to trade outside the laboratory.

Another replicated finding is that the strength of the endowment effect varies with the type of item (Horowitz and McConnell, 2002; Sayman and Öncüler, 2005), such that it is larger for

items that are not normal market goods. Our computational theory of willingness exchange is consistent with this, given the assumption that agents find normal market goods easier to value than goods that are not regularly traded. Greater perceptual error in valuation means greater perceived variability in bids and offers and, hence, a larger endowment effect.

The biggest point of difference between our computational theory of willingness to exchange and previous models of the endowment effect is that the theory depends on not only how agents value goods, but also on their perception of potential bids and offers. This central claim is also consistent with evidence. Experimental subjects with lower WTA (WTP) estimate lower WTP (WTA) among potential trading partners (Van Boven, Dunning and Loewenstein, 2000). Experimental manipulations that affect the distribution of bids and offers alter the WTA-WTP gap. If the value elicitation mechanism is an auction, where bids and offers are irrelevant to the likelihood of exchange, the endowment effect disappears over a few rounds (Shogren et al. 1994). When a uniform price double-auction is employed, which provides direct feedback about the distribution of bids and offers by posting the latest high bid and low offer, the endowment effect is reduced (Franciosi et al., 1996). Weaver and Frederick (2012) show over a series of experiments how the WTA-WTP gap is affected by knowledge of a “market price”. They put this down to changes in utility associated with getting deals near or far from the market price, but our model is also consistent with these effects, because the comparison of own valuation with market price provides a signal of variability in the distribution of bids and/or offers.

Accounting for List’s (2003, 2004) field experiments, in which the endowment effect was absent for experienced dealers at a sports card market, is more problematic. Presumably, valuations of experienced dealers are subject to less perceptual error than those of inexperienced dealers. In our model, however, perceptual error acts via its correlation with the degree of variability in bids and offers, and it is not immediately obvious whether this correlation would be stronger or weaker for experienced dealers. On the other hand, experienced dealers should have higher encounter costs, since List defined experience by the number of routine trades. Higher encounter costs reduce the endowment effect, according to our model. Still, List’s most striking result, and the most difficult to explain, is that experienced sports card dealers did not display any endowment effect in a non-monetary exchange experiment involving mugs and candy bars (List, 2004). According to our model, this result could reflect more accurate perceptions of value, higher encounter costs, and a greater likelihood of understanding the unusual one-shot/one-price market institution. While possible, this explanation is unsatisfactorily general.

Other empirical findings not accounted for by the above model, at least not directly, include those arising from investigations of the potential psychological processes underlying changes in valuation. Various researchers have proposed and tested for particular psychological mechanisms that aim to explain how preferences might be altered by ownership, concentrating primarily on why owning an item might increase its private value. Potential explanations include self-enhancement through motivated taste change (Beggan, 1992), query theory (Johnson, Haubl and Keinan, 2005), focused attention (Ashby, Dickert and Glöckner, 2012) and the salience of attributes (Bordalo, Gennaioli and Shleifer, 2012). Each

account can call on some empirical evidence in support, much of which is recent, and space does not permit a proper review. Yet, from the perspective of the present model, most of these empirical studies assume that changes in valuation measures necessarily imply changes in preferences, i.e. that a change in WTA or WTP over time equates to a change the subject's private value. An alternative is that changes in WTA and WTP over time reflect changes in perceptions of how other people might value the item, since an optimising agent will revise WTA and/or WTP following changes in their perception of potential bids and offers. Future experiments that aim to distinguish between private value and perceptions of the valuations of others might consider collecting data on the "choice equivalent" (CE), defined as the price at subjects are indifferent between money and the item in a binary choice, as a valuation to compare with WTA and/or WTP.

#### *4.2 Novel Predictions*

The computational theory of willingness to exchange makes novel predictions that can be tested with data from the classic WTA-WTP-CE design of Kahneman et al. (1990). Assuming approximate risk neutrality, CE can be identified in our model with the expected value of the agent's private valuation,  $\mu_x$ . Lunn and Lunn (2014) show that this results in a unique prediction regarding variation between subjects in the relationship between CE, WTA and WTP. Subjects with high CE relative to other subjects will have a CE that is closer to WTA, while those with low CE relative to other subjects will have a CE that is closer to WTP. These predictions do not follow from accounts based on loss aversion (Tversky and Kahneman, 1991; Köszegi and Rabin, 2006; Loomes et al., 2009), or those of Kling et al. (2010), Isoni (2011) and Weaver and Frederick (2012).

To test the prediction, although the original data from Kahneman et al. (1990) were unfortunately lost, Lunn and Lunn (2014) obtained original data from Novemsky and Kahneman (2005) and Bateman et al. (2005), both of which closely followed the original design. Across multiple experiments and consumer items, variation in the relativities of CE, WTA and WTP is as predicted by the computational theory of willingness to exchange.

## **5. Conclusion**

Economic literature on the endowment effect tends to centre on whether neoclassical theory or reference dependent theories offer the best account of exchange behaviour. Psychological literature focuses on mechanisms that might lead ownership to change preferences. Both therefore assume that willingness to trade is primarily determined by whether a potential trade increases utility, given the shape of preferences. Our computational theory of exchange represents an alternative, because it focuses not on preferences over outcomes but on the process of exchange itself. The model presented does not require that preferences are stable – ownership can still affect preferences. Nevertheless, it does suggest an alternative explanation of willingness to exchange that does not require a change in preferences to explain the main empirical findings.

The theory combines economic and perceptual theory into a highly generalised model. It shows that under relatively simple and realistic assumptions, where markets involve price



dispersion and sequential encounters, and where value is perceived with significant error, optimal traders will set WTA above WTP, more so the greater the difficulty involved in valuing the good. The model also shows how, in an uncertain environment, comparison of prices ultimately paid and the ease or difficulty of encountering trading partners has heuristic value for agents aiming to optimise WTA and WTP.

The theory can account for the standard experimental findings, is consistent with variability in the strength of the effect with different experimental procedures, and offers an explanation for the increase in endowment effect with nonmarket goods, but does not offer insight into preference formation or change. Its validity as an explanation requires that adaptive behaviour in real markets carries over into laboratory experiments where it is evidently suboptimal. Thus, rather than a laboratory finding that may not occur in the real economy, the endowment effect is modelled as a real world phenomenon that is sometimes absent in the laboratory. Note that this also implies that, while the disparity between WTA and WTP may be a characteristic of behaviour in real markets, the under-trading that occurs in artificial one-shot/one-price laboratory markets may not be a factor in ongoing markets with price dispersion. The model suggests that rather than refusing beneficial trade, agents effectively delay it..

## References

- Ashby, N.J.S., Dickert, S. And Glöckner, A. (2012). Focusing on what you own: biased information uptake due to ownership. *Judgement and Decision Making*, 7, 254-267.
- Baye, M.R., Morgan, J. and Scholten, P. (2006). Information, search and price dispersion, in (Henderschott, T., ed.), *Handbook on Economics and Information Systems*, pp. 326-76, Amsterdam: Elsevier.
- Becker, G.M., DeGroot, M.H. and Marschak, J. (1964). Measuring utility by a single response sequential method. *Behavioural Science*, 9, 226-332.
- Beggan, J.K. (1992). On the social nature of nonsocial perception: the mere ownership effect. *Journal of Personality and Social Psychology*, 62, 229–237.
- Bateman, I., Kahneman, D., Munro, A., Starmer, C. and Sugden, R. (2005). Testing competing models of loss aversion: an adversarial collaboration. *Journal of Public Economics*, 89, 1561-1580.
- Bordalo, P., Gennaioli, N. And Shleifer, A. (2012). Salience in experimental tests of the endowment effect. *The American Economic Review*, 102, 47–52.
- Brodie, E.E. and Ross, H.E. (1984). Sensorimotor mechanisms in weight discrimination. *Perception and Psychophysics*, 5, 477-481.
- Burbeck, C. (1987). Position and spatial frequency in large-scale localization judgements. *Vision Research*, 27, 417-427.
- Ericson, K.M. and Fuster, A. (2013). The endowment effect. National Bureau of Economic Research Working Papers No. 19384.
- Ernst, M.O. and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415, 429-433.

Franciosi, R., Kujal, P., Michelitsch, R., Smith, V. and Deng, G. (1996). Experimental tests of the endowment effect. *Journal of Economic Behaviour and Organisation*, 30, 213-226.

Frederick (2012). Overestimating others' willingness to pay. *Journal of Consumer Research*, 39, 1-21.

Horowitz, J.K. and McConnell, K.E. (2002). A review of WTA/WTP studies. *Journal of Environmental Economics and Management*, 44, 426-447.

Isoni, A. (2011). The willingness-to-accept/willingness-to-pay disparity in repeated markets: loss aversion or 'bad-deal' aversion? *Theory and Decision*, 71, 409-430.

Isoni, A., Loomes, G. and Sugden, R. (2011). The willingness to pay-willingness to accept gap, the "endowment effect," subject misconceptions, and experimental procedures for eliciting valuations: comment. *The American Economic Review*, 101, 991-1011.

Johnson, E.J., Häubl, G. and Keina, A. (2007). Aspects of endowment: a query theory of value construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 33, 461-474.

Kahneman, D., Knetsch, J.L. and Thaler, R.H. (1990). Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy*, 98, 1325-1348.

Kling, C.L., List, J.A. and Zhao, J. (2010). A dynamic explanation of the willingness to pay and willingness to accept disparity. *Economic Inquiry*, 51, 909-921.

Knetsch, J.L. (1989). The endowment effect and evidence of nonreversible indifference curves. *The American Economic Review*, 79, 1277-1284.

Köszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics*, 121, 1133-1166.

List, J.A. (2003). Does market experience eliminate market anomalies? *Quarterly Journal of Economics*, 118, 41-71.

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

Loomes, G., Orr, S. and Sugden, R. (2009). Taste uncertainty and status quo effects in consumer choice, *Journal of Risk and Uncertainty*, 39, 113-35.

Lunn, P.D. and Lunn, M. (2014). What can I get for it? A re-analysis of the endowment effect. ESRI Working Paper No. 385.

Lunn, P.D. and Morgan, M.J. (1997). Discrimination of the spatial derivatives of horizontal binocular disparity, *Journal of the Optical Society of America*, 14, 360-71.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: W.H. Freeman.

Matthews, A. and Stewart, N. (2009). Psychophysics and the judgment of price: judging complex objects on a non-physical dimension elicits sequential effects like those in perceptual tasks. *Judgement and Decision Making*, 4, 64-81.

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

Plott, C.R. and Zeiler, K. (2005). The willingness to pay-willingness to accept gap, the "endowment effect," subject misconceptions, and experimental procedures for eliciting valuations. *The American Economic Review*, 95, 530-545.

Plott, C.R. and Zeiler, K. (2007). Exchange asymmetries incorrectly interpreted as evidence of endowment effect theory and prospect theory? *The American Economic Review*, 97, 1449-1466.

Reinganum, J.F. (1979). A simple model of equilibrium price dispersion. *Journal of Political Economy*, 87, 851-858.

Sayman, S. and Öncüler, A. (2005). Effects of study design characteristics on the WTA-WTP disparity: a meta-analytical framework. *Journal of Economic Psychology*, 26, 289-312.

Shogren, J.F., Shin, S.Y., Hayes, D.J. and Kliebenstein, J.B. (1994). Resolving differences in willingness to pay and willingness to accept. *The American Economic Review*, 84, 255-270.

Tversky, A. and Kahneman, D. (1991). Loss Aversion in Riskless Choice: A Reference-Dependent Model. *The Quarterly Journal of Economics*, 106, 1039-1061.

Van Boven, L., Dunning, D. and Loewenstein, G. (2000). Egocentric empathy gaps between owners and buyers: misperceptions of the endowment effect. *Journal of Personality and Social Psychology*, 79, 66-76.

Weaver, R. and Frederick, S. (2012). A reference price theory of the endowment effect. *Journal of Marketing Research*, 49, 696-707.

## Appendix

### Proof of Proposition 1

Require  $\alpha$  to maximise

$$E(S) = E(Y | Y > \mu_x + \alpha) - \mu_x - \frac{c}{\Pr(Y > \mu_x + \alpha)}.$$

Suppose that  $Y$  has density function  $f(y)$  and cumulative distribution function  $F(y)$ . Then

$$\begin{aligned} E(Y | Y > \mu_x + \alpha) &= \left( \int_{\mu_x + \alpha}^{\infty} y f(y) dy \right) / (1 - F(\mu_x + \alpha)) \\ &= \left( \left[ -y(1 - F(y)) \right]_{\mu_x + \alpha}^{\infty} + \int_{\mu_x + \alpha}^{\infty} (1 - F(y)) dy \right) / (1 - F(\mu_x + \alpha)) \\ &= \mu_x + \alpha + \left( \int_{\mu_x + \alpha}^{\infty} (1 - F(y)) dy \right) / (1 - F(\mu_x + \alpha)). \end{aligned}$$

Hence

$$E(S) = \alpha + \left( \int_{\mu_x + \alpha}^{\infty} (1 - F(y)) dy - c \right) / (1 - F(\mu_x + \alpha)).$$

Differentiate to find  $\alpha^*$

$$1 - (1 - F(\mu_x + \alpha)) / (1 - F(\mu_x + \alpha)) + f(\mu_x + \alpha) \left( \int_{\mu_x + \alpha}^{\infty} (1 - F(y)) dy - c \right) / (1 - F(\mu_x + \alpha))^2 = 0.$$

So  $\alpha^*$  satisfies

$$c = \int_{\mu_x + \alpha^*}^{\infty} (1 - F(y)) dy.$$

Calculation of second order conditions indicates that the second derivative of  $E(S)$  with respect to  $\alpha$  is negative and thus that this is a maximum.

Considering the expression for  $E(S)$  we can see that

$$\max E(S) = \alpha^*$$

(and hence *Proposition 3*) provided that

$$c < \int_{\mu_x}^{\infty} (1 - F(y)) dy.$$

Note we are assuming that  $F(y)$  is continuous and strictly increasing on the range  $(a, b)$  of permitted values for  $Y$  and that  $y(1 - F(y)) \rightarrow 0$  as  $y \rightarrow \infty$ .

**Proof of Proposition 2**

Suppose that  $Y', Y$  have cumulative distribution functions  $F_u, F$  which satisfy the following: for some  $k > 0$  and for all  $y > k$ ,  $F_u(y) < F(y)$ . This implies that  $Y'$  has a fatter upper-tailed distribution than  $Y$ , since  $\Pr(Y > y) = 1 - F(y)$ , similarly for  $Y'$ . Then

$$\int_{y_1}^{\infty} (1 - F_u(y)) dy - \int_{y_1}^{\infty} (1 - F(y)) dy > 0, \text{ whenever } y_1 > k .$$

Thus for solutions  $\alpha^{*'}, \alpha^*$  such that

$$c = \int_{\mu_x + \alpha^{*'}}^{\infty} (1 - F_u(y)) dy = \int_{\mu_x + \alpha^*}^{\infty} (1 - F(y)) dy ,$$

it must hold that  $\alpha^{*'} > \alpha^*$ , whenever  $\mu_x + \alpha^* > k$ .

Note that for normal distributions with the same mean,  $k$  can be taken to be the mean, and furthermore having a fatter upper-tailed distribution is equivalent to having a greater variance.

**Application to normal distribution:**

*Willingness to Accept (WTA)*

Solve for  $\alpha^*$  if

$$\frac{c}{\sigma_y} < \int_{\frac{\mu_x - \mu_y}{\sigma_y}}^{\infty} (1 - \Phi(z)) dz$$

(using transformation to standard normal  $\Phi \sim N(0,1)$  knowing that  $Y \sim N(\mu_y, \sigma_y^2)$ ).

Then to find  $\alpha^*$  we use

$$\frac{c}{\sigma_y} = \frac{1}{\sqrt{2\pi}} e^{-\frac{\lambda^2}{2}} - \lambda(1 - \Phi(\lambda)) \quad , \text{ where } \lambda = \frac{\mu_x - \mu_y + \alpha^*}{\sigma_y} .$$

When this is solved the expected surplus is

$$E(S) = \alpha^* .$$

### Willingness to Pay (WTP)

Suppose amount to pay is  $X \sim N(\mu_x, \sigma_x^2)$ , and that offers are distributed as  $Z \sim N(\mu_z, \sigma_z^2)$ . Suppose WTP is  $\mu_x - \beta^*$ , which satisfies a similar equation, that is  $\beta^*$  maximises expected surplus

$$E(S) = \mu_x - E(Z | Z < \mu_x - \beta) - \frac{c}{\Pr(Z < \mu_x - \beta)}.$$

With similar computations to WTA case we determine  $\beta^*$  according to

$$c = \int_{-\infty}^{\mu_x - \beta^*} F(z) dz,$$

giving  $E(S) = \beta^*$ .

For the normal distribution we get

$$c = \sigma_z \int_{-\infty}^{\eta} \Phi(z) dz$$

so that  $\frac{c}{\sigma_z} = \eta\Phi(\eta) + \frac{e^{-\frac{1}{2}\eta^2}}{\sqrt{2\pi}}$ , where  $\eta = \frac{\mu_x - \mu_z - \beta^*}{\sigma_z}$ .

Year	Number	Title/Author(s) ESRI Authors/Co-authors <i>Italicised</i>
2014	476	Who Feels Inferior? A Test of the Status Anxiety Hypothesis of Social Inequalities in Health <i>Richard Layte</i> and Christopher T. Whelan
	475	Factor Input Substitution in Irish Manufacturing <i>Stefanie A. Haller</i> and <i>Marie Hyland</i>
	474	Changes in Income Distributions and the Role of Tax-benefit Policy During the Great Recession: An International Perspective Olivier Bargain, <i>Tim Callan</i> , Karina Doorley and <i>Claire Keane</i>
	473	International Migration in Ireland, 2012 <i>Philip J. O'Connell</i> and <i>Corona Joyce</i>
	472	Does Bank Market Power Affect SME Financing Constraints? Robert M. Ryan, <i>Conor M. O'Toole</i> , Fergal McCann
	471	Consumption and credit constraints: A model and evidence for Ireland <i>Petra Gerlach-Kristen</i> and <i>Rossana Merola</i>
	470	Socioeconomic Differentials in Male Mortality in Ireland: 1984-2008 <i>Richard Layte</i> and <i>Anne Nolan</i>
	469	Deriving a Method to Estimate Incidence of Stroke in Ireland <i>Maev-Ann Wren</i> and Peter Kelly
	468	The Impact of Parental Income and Education on the Schooling of Children Arnaud Chevalier, Colm Harmon, Vincent O'Sullivan, Ian Walker
	467	SME Credit Constraints and Macroeconomic Effects <i>Petra Gerlach-Kristen</i> , <i>Brian O'Connell</i> and <i>Conor O'Toole</i>
	466	Transitions In and Out of Unemployment Among Young People in the Irish Recession <i>Elish Kelly</i> , <i>Seamus McGuinness</i> , <i>Philip O'Connell</i> , David Haugh and Alberto González Pandiella
	465	The Impact of the Recession on the Structure and Labour Market Success of Young NEET Individuals in Ireland <i>Elish Kelly</i> and <i>Seamus McGuinness</i>

For earlier *Working Papers* see

[http://www.esri.ie/publications/search\\_for\\_a\\_working\\_pape/search\\_results/index.xml](http://www.esri.ie/publications/search_for_a_working_pape/search_results/index.xml)