

Inferential Expectations and the Missing Middle of Price Changes

Timo Henckel

Australian National University^a

Gordon Douglas Menzies

University of Technology, Sydney^b

Daniel John Zizzo

University of East Anglia^c

May 2010

Abstract

Microeconomic evidence suggests price changes are either very small, or large. The theory of inferential expectations predicts this phenomena if agents use a low test size, reflecting a reluctance to change their minds on the basis of evidence.

Keywords: Inferential Expectations, Prices, Near-rationality.

JEL Classification: E31, D84.

^a Research School of Economics, ANU College of Business & Economics, Arndt Building(25), Australian National University, Canberra ACT 0200, Australia (timo.henckel@anu.edu.au).

^b Corresponding author. University of Technology, PO Box 123 Broadway 2007 NSW Australia (gordon.menzies@uts.edu.au). Phone: +61-2-95147728. Fax: +61-2-95147711.

^c School of Economics and CBESS, University of East Anglia, Norwich NR4 7TJ, United Kingdom (d.zizzo@uea.ac.uk).

1. Introduction

Price changes by firms tend to be either very small, or large, creating a ‘missing middle’ in the distribution of price changes (Klenow and Malin, 2008). This note explains this stylized fact from the firm pricing behavior literature by using inferential expectations (IE), a bounded-rational model of belief formation based on hypothesis testing presented and justified in Menzies and Zizzo (2009). Under IE, agents have a null hypothesis (H_0) belief about a variable, which they maintain unless enough evidence exists such that, in a hypothesis test with test size α , they can reject it for an alternative hypothesis (H_1). IE assumes that, when a belief is overturned, agents switch to the rational expectation (RE). Thus, RE is a special case of IE if agents continuously update their beliefs (test size $\alpha=1$). By linking RE to IE, expectations theory gets grounded in the structure of the model, which provides modeling discipline. Properties of IE are: 1) if a rejection region is not reached, agents do not change their mind; 2) a lower test size (α) means greater belief inertia and agents need more extreme evidence to change their mind; 3) agents can make mistakes as they use information.⁴ IE can be considered as a ‘fast and frugal heuristic’ (see Gigerenzer et al., 1999) of belief formation characterized by information-gathering and information-processing costs. It can be related to the idea of sluggish belief adjustment (e.g., Mankiw and Reis, 2002) and to Goldberg and Frydman (1996), who allow agents to conduct hypothesis tests over models. Carroll (2003) provides evidence for sluggish belief adjustment, and Bacchetta and van Wincoop (2009) show how belief adjustment infrequency can be formally related to adjustment cost size. IE might explain exchange rate anomalies such as uncovered interest rate parity failure and delayed overshooting in exchange rates (Menzies and Zizzo, 2009). We do not literally imagine agents to be calculating test statistics and critical values, but these plausible properties make IE an appealing ‘as if’ assumption.

2. The Model

We assume it is optimal for firm i to set the log of price $\ln(p_{it})$ equal to some common economy-wide fundamental e_t , which it cannot observe directly. Nature draws e_t from

⁴ In any hypothesis test 100α per cent of agents mistakenly reject H_0 when H_0 is true.

$N.i.i.d.(0, \sigma_e^2)$ but the firm sees a noisy signal $x_{it} = e_t + v_{it}$ where $v_{it} \sim N(0, \sigma_v^2)$.⁵ All the firms receive a single independent x_{it} signal. As IE agents, they use their own x_{it} for guessing e_t based on a hypothesis test with $N(0, 1)$ critical values $\pm z_{\alpha/2}$. Since e_t is freshly drawn each period, x_{it-1} is irrelevant for guessing e_t . The null hypothesis is rejected if the observed x_{it} is large enough to imply that it is not just measurement error as the null hypothesis would imply:⁶

$$\left. \begin{array}{ll} H_0: e_t=0 \text{ if } & |(x_{it}-0)|/\sigma_v < z_{\alpha/2} \quad \rightarrow \quad \ln(p_{it})=0 \\ H_1: e_t \neq 0 \text{ if } & |(x_{it}-0)|/\sigma_v \geq z_{\alpha/2} \quad \rightarrow \quad \ln(p_{it})=x_{it} \end{array} \right\} \quad (1)$$

Firms believe $\ln(p_{it})=x_{it}$ under H_1 because of the IE assumption that agents make an attempt at a rational calculation at the moment of H_0 rejection (Menzies and Zizzo, 2009).⁷

From (1) a Normal distribution for x_{it} implies a distribution for $\ln(p_{it})$ under the null hypothesis H_0 that is both discrete and continuous (Figure 1). Firms who chose H_0 will set $\ln(p_{it})=0$ creating a discrete probability mass there, while the distribution of $\ln(p_{it})$ for firms who choose H_1 will mimic the continuous density function of x_{it} in the rejection region tails. Assume that the economy has two periods and in each period H_0 is true. Firms receive an i.i.d. x_{it} which assigns them a place in the Figure 1. Table 1 tracks the movement across areas (left tail, middle, right tail) in Figure 1 across the two periods. The proportion $\frac{\alpha}{2}$ of agents in the left tail is divided up across the left, middle and right with proportions $\frac{\alpha}{2} \cdot \frac{\alpha}{2}$, $\frac{\alpha}{2} \cdot (1-\alpha)$ and $\frac{\alpha}{2} \cdot \frac{\alpha}{2}$.⁸

⁵ The assumptions that the mean of e_t is zero and that the variance of v_{it} is known are for simplicity. The central argument would remain intact if each firm received multiple i.i.d. signals and conducted a hypothesis test with an estimated variance.

⁶ We assume that σ_e^2 , the variance of e_t , is not known. Even if it were, the denominator in the test statistic would still be σ_v because conditioning on $e_t=0$ renders it irrelevant. This is no different to an undergraduate test of $H_0: \mu=0$ using a sample mean – one only considers the volatility in the mean, not the volatility of an imaginary prior draw on μ .

⁷ If they knew σ_e^2 they would use the Bayesian signal extraction solution (presented later); this would not affect the key result.

⁸ Readers may convince themselves that Table 1 is sensible for $\alpha=0$ or 1. If $\alpha=0$, there are no rejection regions, so every firm stays in the middle in every period. If $\alpha=1$, the middle vanishes (or, at least, has measure zero) so that one is always in a ‘tail’ (where the two tails now comprise the whole distribution).

3. The Missing Middle

Let us begin by considering the special case where e_t is small and, as a result, Figure 1 is an appropriate approximate description of the outcome of the statistical test. Consider how the area-to-area changes of Table 1 map onto Figure 1. In Figure 2, T_{ij} refers to the terminus of the movement described in Table 1 from the i -th row to the j -th column. For example, T_{12} refers to the movement from the left tail (row 1 in Table 1) to the middle (column 2) which occurs with probability $\frac{\alpha}{2} \cdot (1 - \alpha)$. Double ended arrows reflect that movements occur in both directions.

Belief inertia (a small α) implies that firms staying in the same Figure 2 tail are unlikely to move very much in terms of actual price changes.⁹ Thus, for a small α , there will be some firms with literally zero change (who stay in the middle probability mass of Figure 2, T_{22}), some with small changes (who stay in the same tail in Figure 2, T_{11} and T_{33}), some with large changes (since α is small, T_{21} , T_{12} , T_{32} and T_{23} involve moving a large distance in Figure 2) and some with very large changes (T_{13} and T_{31}). Non-zero price changes smaller than the distance between 0 and the critical values $\pm \sigma_v z_{\alpha/2}$ are the ‘missing middle’ because they are relatively unlikely; movements of this kind would need to take place entirely within the tails, which contain little probability mass for a small α . As α grows, the missing middle vanishes because the tails become so large that ‘staying in the same tail’ no longer implies small price changes, and moving from one tail to another need not imply a very large change. Thus the model explains the missing middle of price changes under the assumption that agents have belief inertia.¹⁰

Let us now generalize our analysis beyond the case where e_t is small. In general, the probability of rejecting the null will depend on the draws of e_t . Will there be a missing middle? To answer this, we need to consider both σ_v^2 and σ_e^2 , since the actual realisations of

⁹ Technically, large changes are possible, but with extremely small probabilities.

¹⁰ In addition, α is a shift parameter for the probabilities of large, and very large, changes (T_{31} , T_{21} , T_{12} , T_{32} , T_{23} and T_{13}). This creates a correlation in the extensive margin: as agents become, say, less belief inert (higher α), the proportion of agents moving prices in *both* directions will increase. For some evidence consistent with this model implication, see Klenow and Krystov (2008).

x_{it} determine whether H_0 is rejected and such actual realisations will depend on $\sigma_e^2 + \sigma_v^2$: specifically, $x_{it} \sim N(0, \sigma_e^2 + \sigma_v^2)$. If σ_v^2 is very small compared with σ_e^2 , then we can show that there is no missing middle. Without loss of generality imagine a fixed σ_e^2 which we normalize to unity. As $\sigma_v^2 \rightarrow 0$ in Figure 1 the ‘accept H_0 ’ region $[-z_{\alpha/2}\sigma_v, z_{\alpha/2}\sigma_v]$ collapses to 0, as per (1). For a fixed σ_e^2 the test statistic will *always* fall in the rejection region and $\ln(p_{it})$ will just be a draw on x_{it} . In this case, the price *change* will simply be the difference between two *N.i.i.d.* draws on x_{it} . If $x_{it} \sim N(0, \sigma_e^2 + \sigma_v^2)$ and $\sigma_v^2 \rightarrow 0$, $x_{it} \sim N(0, \sigma_e^2)$ and $x_{i2} - x_{i1} \sim N(0, 2\sigma_e^2)$, resulting in no missing middle.

Otherwise, the larger σ_v^2 is relative to σ_e^2 , the more the missing middle phenomenon will assert itself. This can be confirmed in computer simulations, two of which are illustrated in Figure 3. They are based on 1000 *i.i.d.* draws of e_t and v_{it} , for $t = 1$ and 2; $\sigma_e = 1$, and the bars show the cases where $\sigma_v = 1/2$ and 1. As Figure 3 shows, a missing middle results.¹¹

4. What about Signal Extraction?

If agents followed RE rather than IE, and knew σ_e^2 , the obvious benchmark would be to engage in Bayesian signal extraction:

$$E(e_t | x_{it}) = \frac{\sigma_v^2 / \sigma_e^2}{1 + \sigma_v^2 / \sigma_e^2} \cdot 0 + \frac{1}{1 + \sigma_v^2 / \sigma_e^2} \cdot x_{it} = \frac{1}{1 + \sigma_v^2 / \sigma_e^2} \cdot x_{it} \quad (2)$$

Since x_{it} is Normal, the *i.i.d.* difference $E(e_2 | x_{i2}) - E(e_1 | x_{i1})$ is Normal too. There would be no ‘missing middle’ in the distribution of changes. There would, of course, also be no missing middle if RE agents did not know σ_e^2 and, as a result, simply chose x_{it} each period.

¹¹ The simulations show that, as agents become less belief inert (α is higher), the proportion of agents moving prices in both directions increases. Thus, the positive correlation in the extensive margin mentioned in an earlier footnote is robust to e_t not being small.

References

Carroll, C.D., 2003. Macroeconomic expectations of households and professional forecasters. *Quarterly Journal of Economics* 118, 269–298.

Klenow, P.J., Kryvtsov, O., 2008. State-dependent or time-dependent pricing: Does it matter for recent U.S. inflation? *Quarterly Journal of Economics* 123, 863-904.

Klenow, P.J., Malin B.A., 2010. Microeconomic evidence on price setting, NBER Working Paper 15826, March.

Menzies, G.D., Zizzo, D.J., 2009. Inferential expectations, *Berkeley Electronic Journal of Macroeconomics (Advances)*, 9.

Gigerenzer, G., Todd, P.M., the ABC Research Group, 1999. *Simple Heuristics That Make Us Smart*. Oxford: Oxford University Press.

Goldberg, M.D., Frydman, R., 1996. Imperfect knowledge and behavior in the foreign exchange market, *Economic Journal* 106, 869-893.

Mankiw, N.G., Reis, R., 2002. Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve, *Quarterly Journal of Economics* 117, 1295-1328.

Bacchetta, P., van Wincoop, E., 2009. Infrequent portfolio decisions: A solution to the forward discount puzzle, *American Economic Review*, forthcoming.

Figure 1. Distribution of prices under H_0




Figure 2. Prices changes


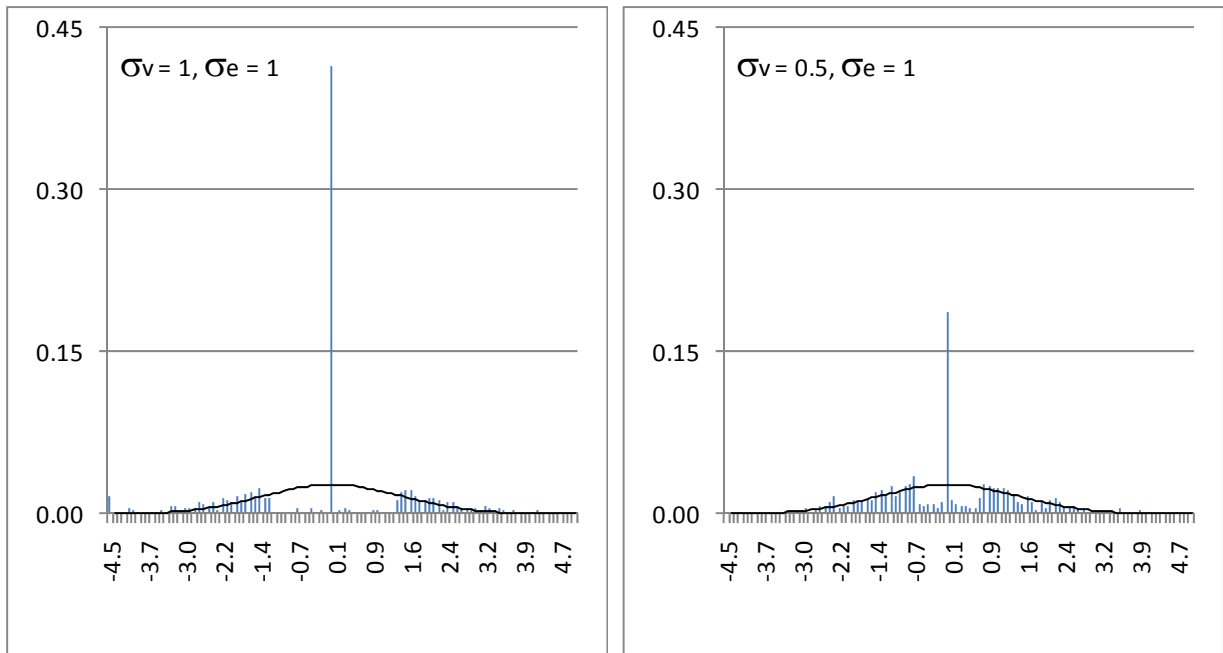


Figure 3. Simulations of price distributions



Note: The Normal distribution is the limiting case where $\sigma_v = 0$ and $x_{i2} - x_{i1} \sim N(0, 2)$.

Table 1. Fraction of firms changing price area from period 1 to period 2

	Left Tail ₂	Middle ₂	Right Tail ₂
Left Tail ₁	$\alpha^2/4$	$(\alpha/2)(1-\alpha)$	$\alpha^2/4$
Middle ₁	$(1-\alpha)(\alpha/2)$	$(1-\alpha)^2$	$(1-\alpha)(\alpha/2)$
Right Tail ₁	$\alpha^2/4$	$(\alpha/2)(1-\alpha)$	$\alpha^2/4$

Notes: The Figure 1 position of the firm's price in period 1 is read off the rows, and that in period 2 off the columns. For example, Left Tail₁ means that the price is in the left tail of Figure 1 in period 1.