

Combining Density and Interval Forecasts: A Modest Proposal*

KENNETH F. WALLIS

*Department of Economics, University of Warwick, Coventry, UK
(e-mail: k.f.wallis@warwick.ac.uk)*

Abstract

The finite mixture distribution is proposed as an appropriate statistical model for a combined density forecast. Its implications for measures of uncertainty and disagreement, and for combining interval forecasts, are described. Related proposals in the literature and applications to the U.S. Survey of Professional Forecasters are discussed.

I. Introduction

Different forecasts of the same event exist whenever forecasters have different information sets or process given information in different ways. The idea that combining such forecasts might be worthwhile rests on the recognition that the combination of forecasts implicitly pools the information sets, and additional information is almost always helpful. The basic idea applies to any kind of forecasts, and has spawned an extensive literature, but this is almost exclusively concerned with point forecasts of random variables. The review and annotated bibliography by Clemen (1989), for example, mostly considers combinations of point forecasts, noting that an issue still deserving attention is the robustness of the simple average of forecasts. Likewise Granger (1989), in

*With apologies to Jonathan Swift (1667–1745). The main ideas in this article were first presented in tutorial lectures at the program on Econometric Forecasting and High-Frequency Data Analysis of the Institute for Mathematical Sciences, National University of Singapore, jointly organized by the School of Economics and Social Sciences, Singapore Management University, May 2004; the kind hospitality of both institutions is gratefully acknowledged. Also acknowledged are the helpful comments of Paul Soderlind and two anonymous referees.

JEL Classification numbers: C53, E37.

his personal commentary 20 years on from the seminal article of Bates and Granger (1969), mostly considers 'developments of the basic idea' insofar as they concern point forecasts. Both articles go beyond this central case to some extent: Clemen considers some event probability forecasting applications; Granger suggests approaches to combining interval and density forecasts. His suggestions have had little impact on subsequent practice, however, and possible reasons for this are discussed below. The general forecast combination literature has maintained its focus on point forecasts, as indicated in several contributions to the compendium edited by Clements and Hendry (2002), and this article aims to widen the focus to include interval and density forecasts.

The relative neglect of combined density forecasts is surprising once it is recalled that they appeared in the original article on the U.S. Survey of Professional Forecasters (SPF), then known as the ASA-NBER Survey (Zarnowitz, 1969), as mean probability distributions of future changes in GNP and prices. Survey respondents are asked not only to report their point forecasts of several variables but also to attach a probability to each of a number of pre-assigned intervals, or bins, into which future GNP growth and inflation might fall. In this way, respondents provide density forecasts of these two variables, in the form of histograms. The probabilities are then averaged over respondents to obtain the mean or combined density forecasts, again in the form of histograms. The reports on the survey results previously published in the *NBER Reporter* and the *American Statistician* did not always refer to the density forecasts, and sometimes combined bins, but mean density forecasts have been included in the press releases of the Federal Reserve Bank of Philadelphia since it assumed responsibility for the survey in 1990 (and changed its name). Initially there was little interest in the individual density forecasts, because of data processing difficulties and variation over time in the number and identity of respondents, although Zarnowitz and Lambros (1987) is a notable exception. More recently, increased interest in density forecasts in general, thanks to publication of such forecasts in several arenas, and the increased accessibility of the SPF in particular, has led to several contributions. Diebold, Tay and Wallis (1999) use the SPF mean density forecast of inflation to illustrate new methods for the evaluation of density forecasts, and Wallis (2003) further extends these methods using the same illustration, among others. The individual SPF responses are used to address several questions of interest by Lahiri, Teigland and Zaporowski (1988), McNees and Fine (1996) and Giordani and Soderlind (2003, 2005). These contributions are further discussed below, following presentation of our preferred statistical framework.

The finite mixture distribution is presented as the appropriate way to think about a combined density forecast in section II, and its implications for analyses of disagreement and uncertainty are discussed in section III. The

alternative statistical framework of Giordani and Soderlind (2003) is considered in section IV, and a simple example in section V illustrates the difference between the two approaches. To combine interval forecasts, it is recommended in section VI that the implied density forecasts be first combined, and then the combined interval forecast with the required probability be read off from the combined density. Some open questions, including a further possible approach to pooling, are described in section VII.

II. A finite mixture distribution

We denote n individual density forecasts of a random variable Y at some future time as $f_i(y)$, $i = 1, \dots, n$. These may come from different forecasters and/or different models and methods, and may be expressed numerically or analytically. For economy of notation time subscripts and references to the information sets on which the forecasts are conditioned are suppressed. The finite mixture distribution is proposed as an appropriate statistical representation for a combined density forecast. It is defined as

$$f_w(y) = \sum_{i=1}^n w_i f_i(y),$$

with weights $w_i \geq 0$, $i = 1, \dots, n$, $\sum w_i = 1$. The same expression appears in statistical decision theory as the linear opinion pool, the commonest form of group consensus probability distribution (see French, 1985, and references therein); again the finite mixture distribution is a relevant statistical model. As these terms suggest, the problem considered in this branch of the decision theory and management science literatures is that of combining expert opinions or subjective probability distributions; Genest and Zidek (1986) provide a further review and annotated bibliography, in which Winkler (1968) is noted as a 'clear forerunner of the Bayesian approach to the aggregation of opinions'. Hitherto, however, there has been little connection to the forecast combination literature.

The finite mixture distribution is well known in the statistical literature, where the central problem is that of identifying and estimating the parameters of the component densities and the mixing proportions. (For a general discussion of finite mixture distributions see Everitt and Hand, 1981, or, for a briefer introduction, Everitt's entry on mixture distributions in Kotz and Johnson, 1985, pp. 559–569, or Stuart and Ord, 1994, Sect. 5.20–5.24.) For a mixture of two normal distributions, this central problem was first considered by Pearson (1894), and mixtures of normal distributions remain a leading case. This case is also relevant to several applications in interval and density forecasting, although the estimation problem is not our present concern.

In reporting probabilities associated with interval forecasts, a normal distribution is often assumed, and some current density forecasts are constructed as normal distributions centred on an associated point forecast and with variance equal to that of past forecast errors. The perspective of a mixture distribution immediately prompts the observation that a combination of such normal density forecasts is not in general a normal distribution, contrary to what is often assumed (by Hendry and Clements, 2004, Sect. 9, for example).

Various weighting schemes appear in the literature on the combination of point forecasts and can be carried over to the present context. Given a series of density forecasts and outcomes weights could be determined from past forecast performance as measured by scoring rules or goodness-of-fit test statistics, or by Bayesian model averaging, for example. However, a surprisingly frequent finding in combining point forecasts is that a simple average, with equal weights, outperforms more complicated weighting schemes. The simple average of the individual histograms has been used to construct the SPF aggregate density forecasts from the beginning, as noted above, and this practice has been followed by several central banks who have subsequently instigated regular surveys of inflation expectations. In the same spirit, we restrict attention in most of what follows to the particular combined density forecast

$$f_C(y) = \frac{1}{n} \sum_{i=1}^n f_i(y). \quad (1)$$

This equally weighted case is sufficiently general to sustain the main points made below, and we return to the general case in the concluding section.

The moments about the origin of $f_C(y)$ are given as the corresponding equally weighted combinations of the moments about the origin of the individual densities. We assume that the individual point forecasts are the means of the individual forecast densities and so denote these means as \hat{y}_i ; the individual variances are σ_i^2 . Then for the first two moments, we have

$$\mu'_1 = \frac{1}{n} \sum_{i=1}^n \hat{y}_i = \hat{y}_C, \quad (2)$$

namely the combined or average point forecast, and

$$\mu'_2 = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i^2 + \sigma_i^2).$$

Hence the variance of f_C is

$$\sigma_C^2 = \mu'_2 - \mu_1'^2 = \frac{1}{n} \sum_{i=1}^n \sigma_i^2 + \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - \hat{y}_C)^2. \quad (3)$$

The first term on the right-hand side of equation (3) is the average individual variance, and the second term is a measure of the dispersion of the individual point forecasts. This decomposition of the variance of the aggregate distribution lies behind several analyses of uncertainty and disagreement, to which we now turn.

III. Consensus and uncertainty

The subtitle echoes the title of the article by Zarnowitz and Lambros (1987), who define 'consensus' as the degree of agreement among point forecasts of the same variable by different forecasters, and 'uncertainty' as the dispersion of the corresponding probability distributions. Their emphasis on the distinction between them was motivated by several previous studies in which high dispersion of point forecasts had been interpreted as indicating high uncertainty. Those studies had not had access to any direct measure of uncertainty, whereas the SPF data provided the opportunity for Zarnowitz and Lambros to check this presumption, among other things. Their definitions are made operational by calculating time series of: (a) the mean of the standard deviations calculated from the individual density forecasts, and (b) the standard deviations of the corresponding sets of point forecasts, for two variables and four forecast horizons. As the strict sense of 'consensus' is unanimous agreement, we prefer to call the second series a measure of disagreement. The two measures are analogous to the two terms on the right-hand side of equation (3), although the use of standard deviations rather than variances breaks the equation; in any event Zarnowitz and Lambros seem unaware of their role in decomposing the variance of the aggregate distribution. They find that the uncertainty (a) series are typically larger and more stable than the disagreement (b) series, thus measures of uncertainty based on the forecast distributions 'should be more dependable'. The two series are positively correlated, however, hence in the absence of direct measures of uncertainty a measure of disagreement among point forecasts may be a useful proxy.

Lahiri, Tiegland and Zaporowski (1988) calculate the first four moments of the individual SPF density forecasts of inflation, and use time series of their average values to examine the effect of inflation uncertainty on the real rate of interest. In passing, they obtain a version of the decomposition in equation (3), although the left-hand side, described as 'total variation in the individual probabilistic forecasts', is not identified as the variance of the aggregate distribution. In their development they work directly with the individual forecast histograms, and there is no consideration of the mean density forecast, nor any underlying statistical model. As with the second moment relationships in equation (3), their average individual skewness and kurtosis

coefficients differ from the skewness and kurtosis coefficients of the aggregate distribution.

IV. Giordani and Soderlind's statistical framework

An alternative statistical model for the SPF mean density forecast is presented by Giordani and Soderlind (2003). They summarize the information set of forecaster i by a scalar signal z_i and write the probability density function of inflation conditional on receiving the signal of forecaster i as $\text{pdf}(\pi|i)$. With π and z_i random variables, the latter having density function $\text{pdf}(i)$, they write the aggregate distribution $\text{pdf}_A(\pi)$ as

$$\text{pdf}_A(\pi) = \int_{-\infty}^{\infty} \text{pdf}(\pi|i)\text{pdf}(i) di. \quad (1')$$

This is Giordani and Soderlind's equation (1), which they say 'amounts to calculating the "marginal" distribution of π ' (quotation marks in original). From the standard relation between the variances of conditional and marginal distributions the variance of the aggregate distribution is then written as

$$\text{Var}_A(\pi) = E(\sigma_i^2) + \text{Var}(\mu_i), \quad (2')$$

which is equation (2) in their article. These equations fulfil similar functions to equations (1) and (3) above, although their statistical foundations are different.

One approach to a mixture distribution is, given a density function dependent on a parameter θ , $f(x|\theta)$, and a weighting distribution for θ , $p(\theta)$ say, then integrating with respect to θ we obtain

$$f(x) = \int_{-\infty}^{\infty} f(x|\theta)p(\theta) d\theta \quad (4)$$

(Stuart and Ord, 1994, p. 181). The marginal distribution $f(x)$ is sometimes known as a compound distribution. One consequence of this approach is that 'any density function $f(x)$ can be viewed as a mixture density simply by imagining extra variables which have been integrated over' (Everitt and Hand, 1981, p. 4). Although equation (4) is at first sight of the form of (1'), it is an inappropriate model for a combined density forecast, because combination in effect pools information sets, rather than integrates them out to obtain a marginal density, which Giordani and Soderlind may be acknowledging when they place marginal in quotation marks. In practice forecasters have diverse information sets, but with much public information in common, and to represent these as a scalar random variable with a well-defined distribution, with aggregation over forecasters then performed by

taking expectations in this distribution, seems unrealistic. We also recall Granger's (1989) remark that 'aggregating forecasts is not the same as aggregating information sets'. The statistical framework in section II is preferred because it directly captures the SPF averaging of individual densities, and to do this it is not necessary to introduce a conditioning random variable. Moreover, the sample average notation on the right-hand side of equation (3) is statistically more accurate than the use of 'E' and 'Var' on the right-hand side of (2').

In the light of their simple model of many forecasters, Giordani and Soderlind find that 'it is not obvious what the aggregate distribution represents'. The point of view of the present article is that it is a combined forecast in the tradition of the point forecast pooling literature. In a second article (Giordani and Soderlind, 2005), however, they show that it is the correct approach to aggregation in an asset-pricing problem in which individual agents have logarithmic utility functions. This last assumption makes aggregation straightforward, and the correspondence does not hold otherwise.

The analysis of the individual SPF inflation forecasts by Giordani and Soderlind (2003) extends the line of research initiated by Zarnowitz and Lambros (1987) and strengthens their conclusion – disagreement is a better proxy for inflation uncertainty than previously thought. Like Zarnowitz and Lambros, Giordani and Soderlind calculate standard deviations, not variances, so identity (3) cannot be checked; unlike them, the standard deviation calculations are based on normal approximations to the forecast histograms. And the time series are now longer, of course. From their evaluation of the individual density forecasts, Giordani and Soderlind conclude that the forecasters underestimated uncertainty. They contrast this finding with that of Diebold *et al.* (1999) who, treating the mean density forecast as that of a representative forecaster, found that it overestimated uncertainty. That such disagreement is possible is obvious from equation (3), and the example of their second article suggests that the choice of measure of collective uncertainty – the variance of the mean density forecast or the average individual variance – will no doubt depend on the problem at hand.

V. Example

A simple univariate example illustrates our statistical framework. Consider the Gaussian AR(2) data generating process

$$Y_t = \phi_1 Y_{t-1} + \phi_2 Y_{t-2} + \varepsilon_t, \varepsilon_t \sim N(0, \sigma_\varepsilon^2),$$

for which the true forecast density of Y_t given observations y_{t-1} and y_{t-2} is

$$Y_t | y_{t-1}, y_{t-2} \sim N(\phi_1 y_{t-1} + \phi_2 y_{t-2}, \sigma_\varepsilon^2).$$

However, two competing forecasters use only a single past observation, lagged one and two periods, respectively, thus their density forecasts are

$$Y_t | y_{t-i} \sim N(\rho_i y_{t-i}, \sigma_i^2), \quad i = 1, 2$$

where ρ_i are autocorrelation coefficients and $\sigma_i^2 = (1 - \rho_i^2)\sigma_y^2$. The combined density forecast is not the marginal distribution of Y_t , as suggested by equation (1') above, but the mixture of normals

$$Y_{Ct} \sim \frac{1}{2}N(\rho_1 y_{t-1}, \sigma_1^2) + \frac{1}{2}N(\rho_2 y_{t-2}, \sigma_2^2).$$

The composite information set for the combined density forecast is identical to the information set of the true forecast density: both contain the same two observations. However the combined forecast uses the information inefficiently, relative to the true forecast density.

In this example, pooling two normal density forecasts produces a non-normal combined forecast, as anticipated above. Its form differs from that assumed for the true density, whereas that of the component forecasts does not. The combined forecast density may be bimodal, if the two past observations are sufficiently different from one another: how different they must be for this to happen depends on the parameter values of the data generating process. This process of aggregating normal densities to produce a non-normal combined forecast is opposite to that observed in the SPF and other surveys, where very disparate individual forecasts typically combine to produce an aggregate forecast with a more normal appearance. Giordani and Soderlind (2005) report that there are several individual histograms in the SPF that look very different from a bell shape, whereas the average distributions are approximately normal. If practical component forecasts assume a specific functional form, circumstances might arise in which it is desired to retain this functional form in the combined forecast, and this possibility is discussed in the final section below. Before that, we retain the present framework to consider the combination of interval forecasts.

VI. Combining interval forecasts

The representation of a combined density forecast as a finite mixture distribution leads to the proposal that interval forecasts be combined via the same route. A density forecast is implicit in an interval forecast, since the calculation of the probability to be attached to an interval requires a distributional assumption, often normality. A combined interval forecast for any required probability can then be obtained from the relevant combined

density forecast, whereas combining intervals directly will not in general give an interval with the correct probability.

The proposal by Granger (1989), applied by Granger, White and Kamstra (1989), attempts to overcome this difficulty by estimating combining weights from data on past forecasts that in effect recalibrate the forecast quantiles. For the forecast cumulative distribution function $F(y)$ define the corresponding quantile function $Q(\theta)$ as

$$Q(\theta) = F^{-1}(\theta), \quad 0 \leq \theta \leq 1.$$

A central interval with forecast coverage p , that is an interval with equal tail probabilities $q/2$, where $q = 1 - p$, is then $(Q(q/2), Q(1 - q/2))$. Given a pair of quantile estimates $Q^A(\theta)$ and $Q^B(\theta)$, Granger suggests the combined estimate

$$Q^C(\theta) = \hat{a}_1(\theta)Q^A(\theta) + \hat{a}_2(\theta)Q^B(\theta).$$

The notation makes it clear that the weights vary with θ , moreover this recalibration requires a historical record of similar forecasts and the corresponding outcomes, which may not always be available. His final proposal is to base a combined density forecast on the distribution function corresponding to $Q^C(\theta)$, presumably evaluated over a grid of values of θ and so requiring several forecast intervals with different coverages – almost a complete density. This seems much more cumbersome than combining the density forecasts directly, as proposed here.

A further example illustrates these issues. To simplify matters we consider two forecast densities differing only in scale, not location, and assume normal distributions. Then the combined density is a mixture of normals with the same mean, which we take to be zero without loss of generality. Both forecasters report 50% and 90% intervals, and we weight the forecasters equally. The density of forecaster 1 is $N(0, 1)$, so the respective intervals are ± 0.67 and ± 1.64 , whereas forecaster 2 is much more uncertain, with forecast density $N(0, 4)$ and associated intervals ± 1.35 and ± 3.29 . With equal weights, combined intervals are then ± 1.01 and ± 2.47 , and in the combined density these have coverage 54% and 88%, respectively, so the directly combined intervals are respectively too wide and too narrow. Intervals with the correct 50% and 90% coverages in the combined density are ± 0.92 and ± 2.62 , respectively, narrower and wider than the directly combined intervals. On seeking unequal weights such that a combination of the individual intervals yields the correct combined interval, we find (0.64, 0.36) for the 50% intervals and (0.41, 0.59) for the 90% intervals: the equally weighted combination of the 50% intervals is too wide, so more weight is needed on the narrower interval of forecaster 1, and vice versa for the 90% interval. These weights vary with the required coverage, as anticipated above, and deriving intervals

from the combined density, rather than seeking to combine individual intervals, is recommended.

VII. Conclusion

Our use of an equally weighted combined forecast to illustrate the statistical framework of the finite mixture distribution leaves on one side the question of the choice of weights in practice. As suggested above, weights could be determined from the track records of the competing forecasts, but several distance measures, scoring rules and goodness-of-fit statistics are available, with little guidance on the choice between them. Compared to the development of the literature on combining point forecasts, the major missing element in density forecast combination is an optimality result comparable to that of Bates and Granger (1969). This shows that a linear combination of two competing point forecasts using the optimal (variance minimizing) weight in general has a smaller forecast mean square error (MSE) than either of the two competing forecasts. It might have seemed counter intuitive that combining with an inferior forecast could improve matters, relative to the MSE measure, but combining using the weight that is optimal with respect to the same measure is seen to achieve this. The only case in which no improvement is possible is that in which one forecast is already the optimal (minimum MSE) forecast; its optimal weight is then 1, and there is no gain in combining with an inferior forecast. Bates and Granger's result is obtained by working analytically in the widely accepted least squares framework, and there is as yet no comparable setting for density forecast combination. One's starting intuition might be the same – that combining a reasonably well-fitting density forecast with one that performs less well makes matters worse – and we wait to see whether this can be overturned as in Bates and Granger (1969).

When individual density forecasts have a common functional form, the finite mixture distribution may not retain this, as seen above, and a situation in which the combined density forecast is required to have the same functional form may need a different combining rule. The problem does not arise for a natural conjugate family of distributions, such as the family of beta distributions, used in Winkler's early examples of the combination of subjective probability distributions, noted above, but these scarcely feature in macroeconomic forecasting. The problem does arise for normal distributions and generalizations such as the three-parameter two-piece normal distribution, used by several central banks to capture an asymmetric balance of risks to their central inflation forecast (Wallis, 1999). The preservation of such a functional form in the combined density forecast may be desired for presentational reasons and/or analytical convenience. As the normal distribution is determined by its first two moments, in such a case pooling could be

achieved via equations (2) and (3), generalized as necessary to include unequal weights, while replacing equation (1) by the specification of the normal distribution for the combined forecast. Comparison of the two combining rules in practice is a subject for future research.

References

- Bates, J. M. and Granger, C. W. J. (1969). 'The combination of forecasts', *Operational Research Quarterly*, Vol. 20, pp. 451–468.
- Clemen, R. T. (1989). 'Combining forecasts: a review and annotated bibliography', *International Journal of Forecasting*, Vol. 5, pp. 559–583.
- Clements, M. P. and Hendry, D. F. (eds) (2002). *A Companion to Economic Forecasting*, Blackwell, Oxford.
- Diebold, F. X., Tay, A. S. and Wallis, K. F. (1999). 'Evaluating density forecasts of inflation: the Survey of Professional Forecasters', in Engle R.F. and White H. (eds), *Cointegration, Causality, and Forecasting: A Festschrift in Honour of Clive W. J. Granger*, Oxford University Press, Oxford, pp. 76–90.
- Everitt, B. S. and Hand, D. J. (1981). *Finite Mixture Distributions*, Chapman and Hall, London.
- French, S. (1985). 'Group consensus probability distributions: a critical survey', in Bernardo J. M., DeGroot M. H., Lindley D. V. and Smith A. F. M. (eds), *Bayesian Statistics 2*, North-Holland, Amsterdam, pp. 183–201.
- Genest, C. and Zidek, J. V. (1986). 'Combining probability distributions: a critique and an annotated bibliography (with discussion)', *Statistical Science*, Vol. 1, pp. 114–148.
- Giordani, P. and Soderlind, P. (2003). 'Inflation forecast uncertainty', *European Economic Review*, Vol. 47, pp. 1037–1059.
- Giordani, P. and Soderlind, P. (2005). 'Is there evidence of pessimism and doubt in subjective distributions? Implications for the equity premium puzzle', *Journal of Economic Dynamics and Control*, in press.
- Granger, C. W. J. (1989). 'Combining forecasts – twenty years later', *Journal of Forecasting*, Vol. 8, pp. 167–173.
- Granger, C. W. J., White, H. and Kamstra, M. (1989). 'Interval forecasting: an analysis based on ARCH-quantile estimators', *Journal of Econometrics*, Vol. 40, pp. 87–96.
- Hendry, D. F. and Clements, M. P. (2004). 'Pooling of forecasts', *Econometrics Journal*, Vol. 7, pp. 1–31.
- Kotz, S. and Johnson, N. L. (eds) (1985). *Encyclopedia of Statistical Sciences*, Vol. 5, Wiley, New York.
- Lahiri, K., Teigland, C. and Zaporowski, M. (1988). 'Interest rates and the subjective probability distribution of inflation forecasts', *Journal of Money, Credit, and Banking*, Vol. 20, pp. 233–248.
- McNees, S. K. and Fine, L. K. (1996). *Forecast Uncertainty: Can it be Measured?* Paper presented at the Conference on Expectations in Economics, Federal Reserve Bank of Philadelphia, October 1996.
- Pearson, K. (1894). 'Contributions to the mathematical theory of evolution', *Philosophical Transactions of the Royal Society of London*, A, Vol. 185, pp. 71–110.
- Stuart, A. and Ord, J. K. (1994). *Kendall's Advanced Theory of Statistics*, 6th edn, Vol. 1, Edward Arnold, London.
- Wallis, K. F. (1999). 'Asymmetric density forecasts of inflation and the Bank of England's fan chart', *National Institute Economic Review*, No. 167, pp. 106–112.

- Wallis, K. F. (2003). 'Chi-squared tests of interval and density forecasts, and the Bank of England's fan charts', *International Journal of Forecasting*, Vol. 19, pp. 165–175.
- Winkler, R. L. (1968). 'The consensus of subjective probability distributions', *Management Science*, Vol. 15, pp. B61–B75.
- Zarnowitz, V. (1969). 'The new ASA-NBER survey of forecasts by economic statisticians', *American Statistician*, Vol. 23, pp. 12–16.
- Zarnowitz, V. and Lambros, L. A. (1987). 'Consensus and uncertainty in economic prediction', *Journal of Political Economy*, Vol. 95, pp. 591–621.