International Review of Economics and Finance 45 (2016) 453–469

Contents lists available at ScienceDirect

International Review of Economics and Finance

journal homepage: www.elsevier.com/locate/iref

# Option-implied probability distributions: How reliable? How jagged?☆

Marco Taboga

*Economic Outlook and Monetary Policy Department, Bank of Italy,Via Nazionale 91, 00184 Roma, Italy*

ARTICLE INFO

*Article history:*
Received 7 January 2016
Received in revised form 20 July 2016
Accepted 21 July 2016
Available online 25 July 2016

*JEL classification:*
C14
C58
G13

*Keywords:*
Implied state prices
Implied risk-neutral distributions
State price estimation
Option-implied distributions

ABSTRACT

Estimates of option-implied probability distributions are routinely used in central banks, as well as in other institutions, but their reliability is often difficult to assess. To address this issue, we propose a semi-nonparametric model that allows to compute exact credible intervals around estimated distributions. By analyzing a panel of S&P 500 options, we find that the estimates of the distributions are quite precise. We also provide evidence that the multi-modality often found in option-implied distributions could be an artifact due to over-fitting, and that models with uni-modality constraints have high posterior odds.


## 1. Introduction

Since the seminal contributions of Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978), the problem of estimating risk-neutral probabilities (or, equivalently, state prices[1] ) from cross-sections of option prices has received much attention in the literature. Furthermore, several papers have demonstrated the importance of estimating state prices for various empirical applications, including the estimation of investors' risk preferences, the assessment of market beliefs about events of interest, the pricing of exotic derivatives, the estimation of parametric asset pricing models, and the calibration of risk management models (see Bondarenko, 2003 for a review of applications).

We propose a new model to estimate state prices, which belongs to the class of semi-nonparametric models (see Gallant, 1987, and Fengler & Hin, 2015). In our model, the vector of state prices is assumed to be interpolated by a spline function.

☆ All the views expressed in this paper are mine and do not necessarily reflect those of the Bank of Italy. I thank Giuseppe Grande, Marcello Pericoli and an anonymous referee for helpful comments and suggestions.
*E-mail address:* marco.taboga@bancaditalia.it.

[1] State prices are the prices of Arrow-Debreu securities, that is, of securities that pay one unit of account in a specific state of nature and nothing in any other state. They coincide with risk-neutral probabilities when the risk-free interest rate is equal to zero. Otherwise, when the risk-free rate is different from zero, risk-neutral probabilities are proportional to state prices.

http://dx.doi.org/10.1016/j.iref.2016.07.013
1059-0560/

We show that this assumption is equivalent to imposing a set of linear equality restrictions on state prices (a similar result is found by Monteiro, Tütüncü, and Vicente, 2008, who show that option prices are linear in spline parameters). We propose two estimators of the vector of state prices. The first one is a least absolute deviations (LAD) estimator that can be obtained as the solution of a computationally inexpensive linear programming problem. The second one is a Bayesian estimator. We prove that the Bayesian estimation problem boils down to updating the prior on the coefficients of a linear regression equation.

We use our model to tackle two issues that, to our knowledge, have been seldom addressed in the literature on the nonparametric estimation of state prices.

The first issue concerns the uncertainty that surrounds the estimates of state prices. While some of the existing models allow to derive the asymptotic distributions of state price estimators, asymptotics are a poor guide for predicting the behavior of the estimators because of the small sample sizes that typically characterize the cross sections of option prices (Aït-Sahalia & Duarte, 2003). Our Bayesian estimation procedure instead allows to derive exact credible intervals around estimated state prices.

We believe that computing credible intervals precisely is important not only from a statistical viewpoint, but also in view of the economic applications of state price estimation. Estimated state prices and risk-neutral densities are used to infer market beliefs about economic events of interest, for instance, in central banks, where they are routinely used for monetary policy purposes (e.g., Malz, 2014; Söderlind & Svensson, 1997). To understand how much weight these estimates should have in influencing policy decisions, it is essential to assess their reliability and statistical precision.

Furthermore, the statistical precision of the estimates has significant consequences for the pricing of derivatives. For example, suppose that a dealer bank uses estimated state prices to price bespoke derivatives traded with clients. It can be rational (e.g., Cont, 2006; Routledge & Zin, 2009) for the bank to adjust the spread between bid and ask quotes so as to seek protection from the uncertainty about the fair value of the bespoke derivative. Because this uncertainty stems from the uncertainty about the true value of the state prices, the bid and ask quotes ultimately depend on the statistical precision of the estimates of the state prices. Furthermore, in the absence of a proper quantification of the precision of the estimates, the fair value of the bespoke derivative is uncertain in the Knightian sense of the term, and the uncertainty can give rise to ambiguity (e.g., Ellsberg, 1961), with the well-known behavioral implications in terms of pricing (e.g., Dana, 2004; Easley & O'Hara, 2010).

By analyzing a panel of options on the S&P 500 stock index with our model, we find that the posterior distributions of the state prices and of the risk-neutral probabilities have low dispersion (the standard deviation of the state prices is on average 1.7% of their value). The posterior dispersion is even lower for the quantiles of the risk-neutral probability distributions. In particular, while it has been conjectured (e.g., Lee, 2014) that tail quantiles might be difficult to estimate precisely with nonparametric methods, we find that the credible intervals around the estimated tail quantiles are quite tight. According to some experiments we have conducted, this finding might be explained by the fact that information about tail quantiles is provided not only by illiquid deep-out-of-the-money options (say, for lower quantiles, the put options whose strike is close to the quantile) but also by liquid in-the-money-options.

The second issue we address is the multi-modality of the risk-neutral probability distributions derived from estimated state prices.[2] Estimated risk-neutral distributions are often reported to be multi-modal (e.g., Fengler & Hin, 2015; Melick & Thomas, 1997; Taylor, Tzeng, & Widdicks, 2012). Some studies have advanced possible explanations for multi-modality: heterogeneity in investors' beliefs (Ziegler, 2002), jumps and correlation between volatility and returns (Frachot, Laurent, & Pichot, 1999), regime-switches (see Chabi-Yo, Garcia, & Renault, 2008 for a review). Despite the attempts to rationalize it, multi-modality is still deemed economically implausible by many economists, on the grounds of a monotonicity argument (the more extreme an outcome is, the less likely it should be). Furthermore, as suggested by Malz (2014) and Fengler & Hin (2015), multi-modality could be authentic, but it could also be an artifact of the techniques used to estimate risk-neutral distributions. Given the diversity of views on this issue, it is desirable to have models that allow to impose the prior of uni-modality and to see what impact it has on estimated risk-neutral distributions and option pricing functions. For example, the argument in favor of uni-modality could be stronger if the prior had a minor impact on pricing errors, that is, on the ability to accurately fit option prices. To our knowledge, how to impose uni-modality restrictions within a nonparametric option pricing framework has been seldom discussed in the literature (although Fengler & Hin, 2015, in suggesting directions for future research, hint to a possible extension of their model that could accommodate uni-modality). For this reason, we thoroughly discuss how to incorporate the prior of uni-modality in the Bayesian estimation of our model and what consequences the prior has.

We find that the risk-neutral probability distributions estimated from our data are often multi-modal. Imposing the prior of uni-modality can significantly change the shape of the estimated risk-neutral densities, but it has limited effects on the pricing errors and on the estimated quantiles of the risk-neutral distribution of returns. Furthermore, by computing posterior odds ratios between uni-modal and multi-modal models, we find that uni-modality is, a posteriori, much more likely than multi-modality. We interpret our findings as empirical support for the hypothesis that the multi-modality often found in option-implied distributions could be an artifact due to over-fitting.[3]

As reported by Lee (2014), the smoothed implied volatility smile (SML) method, first proposed by Shimko (1993), is one of the most widely used techniques for estimating risk-neutral densities and state prices. The SML method, in its several variants (e.g.,

[2] While in this paper we propose estimators of state prices, our estimators can be easily rescaled to produce estimators of risk-neutral probability distributions.

[3] In any case, the ability to compute odds ratios between multi-modal and uni-modal models should provide a means of discriminating between true and spurious multi-modality.