Market-Based Expectations as a Tool for Policymakers*

March 31, 2016

Abstract

The paper considers the decision problem of a policymaker who is uncertain about the net benefits of an action choice. We compare the equilibrium outcomes for two different policymaker objectives. The first is a representative household's ex-ante utility. The second is the market-based expectation of net benefits imputed from financial asset prices. We show that, if asset markets are complete, the equilibrium outcomes are the same for two objective functions. Equilibrium outcomes are typically different if the policymaker instead maximizes the statistical expectation of next benefits. We illustrate our results using the example of an inflation-targeting central bank.

*Authors: Ron Feldman (FRB-Minneapolis), Kenneth Heinecke (FRB-Minneapolis), Narayana Kocherlakota (University of Rochester), Samuel Schulhofer-Wohl (FRB-Minneapolis), Thomas Tallarini (FRB-Minneapolis). This version of the paper builds on an earlier (January 2015) draft called "Market-Based Probabilities as a Tool for Policymakers." The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Please address any comments to nkocherl@ur.rochester.edu.

1 Introduction

Policymakers often make decisions in the context of uncertainty about the net benefits of their action choice. As a result, in making a choice, the policymaker has to use an objective function that trades off net benefits across the various states of the world. In this paper, we consider three possible objectives that policymaker can use to make these cross-state trade-offs. The first objective is to maximize the ex-ante utility function of a typical household. The second is to maximize expected net benefits (where the expectation is based on the policymaker's assessment of the likelihood of the various outcomes). The final objective is to maximize the *market-based expectation* of net benefits. Here, the market-based expectation of a particular random variable X is the price of an asset that pays off X in terms of a risk-free asset.[1]

Our main result is that, under certain conditions, the policymaker who maximizes the market-based expectation of net benefits achieves the same outcome as the policymaker who maximizes household ex-ante utility. In contrast, the policymaker who maximizes the statistical expectation of net benefits will typically make a different choice. The intuition behind these results is simple. Because of risk aversion, households typically assign a higher marginal value to resources in states of the world associated with adverse economic outcomes (like a recession) than resources in other possible states with better economic outcomes. Objective probabilities, by their very construction, do not embed this aspect of household preferences. Financial market prices do embed this aspect of household preferences - and so market-based

[1]Market-based probabilities are often termed risk-neutral probabilities in financial economics.

probabilities, by their very construction, also do. It follows that policymakers who want to be reflective of this important aspect of household preferences should be guided by market-based probabilities rather than by objective probabilities.

Others before us have discussed how policymakers can find financial market prices to be a useful guide to decision-making.[2] However, our emphasis is quite distinct from this prior literature. The earlier papers saw financial market prices as potentially useful as a source of information about underlying true or objective probabilities of possible futures. In contrast, we emphasize that market-based probabilities imputed from financial market data will typically be quite different from objective probabilities estimated using statistical models. We see market-based expectations as being useful exactly because of this difference, which is informative about the degree of household's aversion to losses generated by policymaker's decisions.

Others have argued that financial market prices are *not* a useful guide to policymakers. For example, in the context of inflation-targeting, Bernanke and Woodford (1997) say that "asset price measures of expected inflation are likely to be contaminated by ... changes in the inflation risk premium." In our model, this purported "contamination" is due to cross-state variation in the the representative household's aversion to the losses generated by the policymaker's decision. What Bernanke and Woodford see as "contamination" is exactly why policymakers should maximize the market-based, as opposed to statistical, expectation of net benefits.

One way to view our main result is that it extends the basic principles of *intertemporal* policy choice to policy choice under *uncertainty.* Economists generally

[2]See Hetzel (1992) and Sumner (1995) for example.

agree policymakers should use one kind of financial market prices - interest rates - as a benchmark approach to weighting resources at different points in time. Our main result extends this perspective by showing that policymakers can maximize social welfare by using another aspect of financial market prices – market-based probabilities - as a benchmark weighting of resources in different possible outcomes.[3]

Section 2 presents a simple abstract model. Section 3 discusses the main results. Section 4 illustrates the results in the context of an inflation-targeting example. Section 5 concludes.

2 Three Policy Games

In this section, we examine a simple abstract model in which a policymaker makes a decision under uncertainty. We consider three policy games that are distinguished by the objective of the policymaker. We first describe the common elements of the three games and then describe the different policymaker objectives.

2.1 Common Elements

In all of the games, there are households and a policymaker and there are three periods: a *trading period*, a *planning period*, and a *realization period.* During the *trading period*, households trade (a complete set of) financial securities. Then, during the *planning period*, the policymaker chooses an action, a, today that affects outcomes

[3] The Federal Reserve Bank of Minneapolis reports the market-based probabilities of various events on its website, including changes in inflation, interest rates and other asset values. The website — which offers users the option of receiving updated data and commentary — can be found at https://www.minneapolisfed.org/banking/mpd.

in the *realization period*. In addition, the outcome of the action in the realization period depends on the realization of a random variable x, with N possible realizations $\{x_n\}_{n=1}^{N}$.

The action chosen by the policymaker has costs and benefits that depend on the realization of the state variable, x. Let $B(a,x)$ denote the net benefits (gross benefits minus costs) associated with action a in state x. Since $B(a,x)$ measures *net* benefits, its realization may be positive or negative. For all x, the function B is strictly concave as a function of a.[4]

In the trading period, households are identical. They are each endowed with $\bar{y}$ units of consumption, where $\bar{y}$ is the same across states. Within a given state n, a household's utility is given by:

$$U_n = u(c(x_n)) + B(a, x_n) \tag{1}$$

where u is strictly increasing and concave and $c(x_n)$ represents consumption in state n. Households' ex-ante utility (before they know which state occurs) is given by:

$$V(U_1, ..., U_N) \tag{2}$$

[4]This structure assumes that the net benefit function B depends on the policymaker's action a and random influences x that are wholly independent of a. This restriction is without loss of generality when x is continuous. In particular, suppose that B is a function of (z, a, x), where F is the cumulative distribution function of z, conditional on x and a. In this formulation, B depends on some random influence z that is influenced by the policymaker's choices. However, we can create an entirely isomorphic model by defining $\widehat{B}(u, a, x)$, where u is uniform $[0, 1]$, to be equal to $B(F^{-1}(u|a, x), a, x)$. (This is the same trick that underlies most Monte Carlo simulation experiments.) In this isomorphic model, the random influences on $\widehat{B}$ are independent of the policymaker's choice a.

where V is strictly increasing and strictly concave.[5] We assume too that Vsatisfies standard Inada conditions (so that the partial derivative V_n converges to infinity (zero) as U_n converges to zero (infinity)).

We assume that the background household endowment is constant across states. More generally, we can allow the household's endowment to be random, but we then have to change (1) to be:

$$U_n = u(c(x_n) + B(a, x_n)) \tag{3}$$

Note that the within-state marginal rate of substitution between consumption and net benefits is constant across states in both specifications (1) (with a constant endowment) and (3) (with a potentially random endowment). This restriction that the within-state marginal rate of substitution between consumption goods and the net benefits is constant over states ensures that there is a meaningful sense in which the net marginal benefit of the policymaker's action is in units of consumption. It is essential for our results.[6]

2.2 Social Welfare Game

The policy games differ in terms of the objective function of the policymaker in the planning period. In the *social welfare* game, the policymaker chooses a so as to

[5] As Armenter (2015) points out, the assumption of strict concavity is crucial for our results.

[6] We use (1) because it allows us to embed Woodford's (2003) assessment of the welfare costs of inflation into our analysis in Section 4.

maximize social welfare:

$$V((u(\bar{y}) + B(a, x_n))_{n=1}^N)$$

Note that this formulation of the objective function assumes that, as will be true in equilibrium, the identical households do not trade in the asset market. The unique equilibrium[7] a^*_{SWF} of this game is characterized by the first order condition:

$$\sum_{n=1}^{N} V_n(a^*_{SWF}) B_a(a^*_{SWF}, x_n) = 0$$

Here, B_a represents the partial derivative of B with respect to its first argument, and $V_n(a)$ represents the partial derivative of V with respect to its nth argument given the policymaker chooses action a:

$$V_n(a) \equiv \frac{\partial V}{\partial U_n}((u(\bar{y}) + B(a, x_n))_{n=1}^N)$$

conditional on the policymaker's choosing a.

2.3 Statistical Expectation Game

In the statistical expectation game, the policymaker maximizes expected net benefits. Formally, let $(p_1, ..., p_N)$ be the policymaker's assessment of the likelihood of the

[7] We can establish uniqueness as follows. Define the function $\Gamma(a) = V((u(\bar{y}) + B(a, x_n))_{n=1}^N)$. Suppose $a_\lambda = \lambda a + (1-\lambda)a'$. Then, $B(a_\lambda) > \lambda B(a) + (1-\lambda)B(a')$ because B is strictly concave. Since V is strictly increasing, we can conclude that $\Gamma(a_\lambda) > V((u(\bar{y}) + \lambda B(a, x_n) + (1-\lambda)B(a', x_n))_{n=1}^N) > \lambda\Gamma(a) + (1-\lambda)\Gamma(a')$, where the last inequality follows from the strict concavity of V. Hence, Γ is strictly concave, and there is a unique solution to $\Gamma'(a^*_{SWF}) = 0$.

various states. Then, the policymaker chooses a so as to maximize:

$$\sum_{n=1}^{N} p_n B(a, x_n)$$

The equilibrium a^*_{SE} to this game is characterized as the unique solution to the equation:

$$\sum_{n=1}^{N} p_n B_a(a^*_{SE}, x_n) = 0$$

2.4 Market-Based Game

In the *market-based* game, policymakers maximize the market-based expectation of net social benefits.

Formally, the policymaker observes the outcome of household interactions in the trading period. Let $q_n(\widehat{a})$ denote the implied price today of goods in state n, conditional on households' common beliefs $\widehat{a}$ in the trading period about the policymaker's action choice in the planning period. Now, define

$$q^*_n(\widehat{a}) = \frac{q_n(\widehat{a})}{\sum_{n=1}^{N} q_n(\widehat{a})}. \tag{4}$$

Since $q_n(\widehat{a})$ is the price of goods in state n, $q_n(\widehat{a}) \geq 0$ for all n. As a result, $q^*_n(\widehat{a}) \geq 0$ for all n. In addition, $\sum_{n=1}^{N} q^*_n(\widehat{a}) = 1$. Therefore, $\{q^*_n(\widehat{a})\}_{n=1}^{N}$ is a probability measure over the states of the world. We will call this the *market-based* probability measure.[8] Given the market-based probability measure, we can define a new expec-

[8]The vector $\{q^*_n\}$ is often referred to as a *risk-neutral* probability measure, especially in finance.

tations operator, E^*, over any random variable ϕ:

$$E^*[\phi|\widehat{a}] = \sum_{n=1}^{N} q_n^*(\widehat{a})\phi_n, \tag{5}$$

where ϕ_n is the realization of the random variable ϕ in state x_n.

Given this definition, in the market-based game, the policymaker's objective function in the planning period, conditional on the households' common belief $\widehat{a}$, is given by:

$$E^*[B(a,x)|\widehat{a}]$$

We can then solve for the equilibrium of the market-based game using backwards induction. Conditional on household beliefs $\widehat{a}$, the solution to the policymaker's problem in the planning period is characterized by the first order condition:

$$\sum_{n=1}^{N} q_n^*(\widehat{a})\frac{\partial B}{\partial a}(a_{MKT}^*, x_n) = 0$$

The market-based expectation q^* is related to the marginal utility of consumption of the various households:

$$q_n^*(\widehat{a}) = \frac{V_n(\widehat{a})u'(\bar{y})}{\sum_{n=1}^{N} V_n(\widehat{a})u'(\bar{y})} = \frac{V_n(\widehat{a})}{\sum_{n=1}^{N} V_n(\widehat{a})}.$$

In equilibrium, the households' beliefs about the policymaker's choice are correct. The equilibrium a_{MKT}^* can be characterized as the solution to the equation:

$$\frac{\sum_{n=1}^{N} V_n(a_{MKT}^*)B_a(a_{MKT}^*, x_n)}{\sum_{m=1}^{N} V_m(a_{MKT}^*)} = 0$$

3 Results

In this section, we establish an equivalence result between the equilibria in the social welfare game and in the market-based game. We also show that the equilibrium in the statistical expectation game is typically different, in a way that is reflected in the magnitude of the risk premium on a particular asset.

3.1 Equivalence Result

It is straightforward to prove that the equilibrium outcomes of the social welfare game and the market-based expectation game are identical.

Theorem 1. *The policymaker's equilibrium action a^*_{SWF} in the social welfare game is the same as the policymaker's equilibrium action a^*_{MKT} in the market-based expectation game.*

Proof. The equilibrium in the social welfare game is the unique solution to the equation:

$$\sum_{n=1}^{N} V_n(a^*_{SWF}) B_a(a^*_{SWF}, x_n) = 0$$

The equilibrium in the market-based game is the unique solution to the equation:

$$\frac{\sum_{n=1}^{N} V_n(a^*_{MKT}) B_a(a^*_{MKT}, x_n)}{\sum_{m=1}^{N} V_m(a^*_{MKT})} = 0$$

The two equations are identical, modulo the positive denominator in the latter equation. $\square$

Intuitively, financial market prices reflect households' willingness to substitute resources across states. Hence, maximizing the market-based expectation trades off resourcs across states just as households would.

3.2 Non-Equivalence Result 1

In this subsection, we consider the non-equivalence of the equilibrium outcomes of the social welfare game and the statistical expectation game. Of course, at this point, there is no connection between the households' ex-ante utility function V and the policymaker's assessment $(p_1, ..., p_N)$ of the likelihoods of the various states of the world. Without such a connection, there is no reason to expect the outcomes in the two games to align. And, in our view, this lack of connection (between households' ex-ante utility and the policymaker's forecasts) seems like a plausible description of reality.

However, the possibility of risk aversion means that even if households do form expectations in the same way that the policymaker does, the outcome in the statistical expectation game is generally not the same as that in the social welfare game. In particular, suppose V takes the form:

$$V(U_1, ...U_N) = \sum_{n=1}^{N} p_n \phi(U_n) \tag{6}$$

The households maximize expected utility. Their cardinal utility function is the composition of ϕ and the within-state utility function u. Their expectation is formed using the policymaker's assessment of the likelihoods of the various states.

Theorem 2. *Suppose V satisfies (6), where ϕ is strictly concave. Suppose that $B_a(a,x) = \xi(B(a,x))$, where ξ is strictly decreasing, and that there exists states n,m such that $B(a^*_{SE}, x_n) \neq B(a^*_{SE}, x_m)$. Then, a^*_{SE} is not equal to a^*_{SWF}.*

Proof. We proceed by contradiction. Suppose that $a^*_{SE} = a^*_{SWF}$. Then:

$$\begin{aligned}
0 &= \sum_{n=1}^{N} p_n \phi'(u(\overline{y}) + B(a^*_{SE}, x_n)) B_a(a^*_{SE}, x_n) \\
&= \sum_{n=1}^{N} p_n [\phi'(u(\overline{y}) + B(a^*_{SE}, x_n)) - \sum_{m=1}^{N} p_m \phi'(u(\overline{y}) + B(a^*_{SE}, x_n))] B_a(a^*_{SE}, x_n) \\
&= \sum_{n=1}^{N} p_n [\phi'(u(\overline{y}) + B(a^*_{SE}, x_n)) - \sum_{m=1}^{N} p_m \phi'(u(\overline{y}) + B(a^*_{SE}, x_n))] \xi(B(a^*_{SE}, x_n))
\end{aligned}$$

This last difference is positive, because, over the various states, $\xi(B(a^*_{SE}, x_n))$ is strictly increasing as a function of $[\phi'(u(\overline{y}) + B(a^*_{SE}, x_n)) - \sum_{m=1}^{N} p_m \phi'(u(\overline{y}) + B(a^*_{SE}, x_n))]$. That's a contradiction. □

By using the statistical expectation, the policymaker ignores the households' risk aversion (encoded in ϕ). By doing so, the policymaker makes a different choice than is obtained by maximizing social welfare.

3.3 Non-Equivalence Result 2

In this subsection, we show that the existence of risk premia is a sign that the policymaker should not use an objective based on statistical expectations.

Consider an asset with state-contingent payoff $(\theta_n)_{n=1}^{N}$. Suppose households were risk-neutral and had the same assessment about the likelihood of various states as the policymaker. Then the price of this asset in terms of risk-free bonds would be:

$$(\sum_{n=1}^{N} \theta_n p_n)$$

Suppose households believe that the policymaker will play a^*_{MKT}. The actual equilibrium price of the asset, in terms of risk-free bonds, is given by:

$$\sum_{n=1}^{N} \theta_n q^*_n(a^*_{MKT})$$

We define the *risk premium* of the asset to be the deflation in its price attributable to households' risk aversion:

$$RP(\theta) = \sum_{n=1}^{N} \theta_n [p_n - q^*_n(a^*_{MKT})]$$

Theorem 3. *Consider an asset that pays off the marginal benefit in each state:*

$$MB = (B_a(a^*_{MKT}, x_n)_{n=1}^{N})$$

Then its risk premium $RP(MB) = 0$ if and only if $a^*_{MKT} = a^*_{SE}$.

Proof. Definitionally, the price of the asset is zero, because:

$$\sum_{n=1}^{N} B_a(a^*_{MKT}) q^*_n(a^*_{MKT}) = 0$$

If the risk premium is zero, then:

$$\sum_{n=1}^{N} B_a(a^*_{MKT})p_n = 0$$

and $a^*_{MKT} = a^*_{SE}$. Conversely, if $a^*_{MKT} = a^*_{SE}$, then $RP(MB) = 0$. $\square$

The existence of a non-zero risk premium is a sign that households are averse to risk. In that case, the policymaker should not use the objective function based on statistical expectations exactly because it ignores households' risk aversion.[9]

4 Example: Inflation-Targeting

In this section, we illustrate our results using the example of an inflation-targeting central bank. In this example, the benefit function B is given by:

$$B(a, x) = -\gamma(\pi(a, x) - \pi^*)^2$$

Woodford (2003, p. 399) shows that this quadratic function approximates the loss due to relative price distortions generated by Calvo pricing frictions. (One inessential difference is that the above approximate loss function is derived under the assumption of indexation, so that the non-optimizing price-setters increase their prices at rate π^*; Woodford sets $\pi^* = 0$.)

We assume as well that the realized level of inflation is given by a linear combi-

[9]Cochrane (2011) argues that risk premia on many assets are highly variable.

nation of the central bank's choice a and a disturbance term x:

$$\pi(a, x) = a + x$$

The disturbance term x has N possible realizations $(x_n)_{n=1}^N$, and has mean zero according to the policymaker's assessment $(p_n)_{n=1}^N$ of their likelihoods.

Finally, we assume that households' ex-ante utility over random inflation outcomes is given by:

$$\sum_{n=1}^{N} \widehat{p}_n \phi(u(\bar{y}) + B(a, x_n), x_n)$$

where ϕ is strictly increasing and weakly concave. This formulation of household preferences differs from that of Woodford (2003), because he restricts ϕ to be linear and assumes that the policymaker and the representative household use the same probability density to form their expectations. The assumption of linearity in his formulation imposes a tight restriction between the households' willingness to substitute consumption across states (risk aversion) and their intra-state willingness to substitute consumption for lower relative price distortions. This restriction is inconsistent with many features of the asset pricing data (including the behavior of inflation risk premia). Our formulation relaxes that restriction in a way that is directionally consistent with these data.

In this example, we can use the quadratic nature of B to characterize the equilibrium policy choices in the three games:

$$\begin{aligned}
\sum_{n=1}^{N} \widehat{p}_n \phi'(u(\bar{y}) - \gamma(a_{SWF} + x - \pi^*)^2, x_n)\gamma(a_{SWF} + x - \pi^*) &= 0 \\
\sum_{n=1}^{N} p_n \gamma(a_{SE} + x - \pi^*) &= 0 \\
\sum_{n=1}^{N} \widehat{p}_n \phi'(u(\bar{y}) + \gamma(a_{SWF} + x - \pi^*)^2, x_n)\gamma(a_{MKT} + x - \pi^*) &= 0
\end{aligned}$$

These solutions can be characterized in a more familiar way as follows:

$$\begin{aligned}
E^*(\pi(a^*_{SWF}, x)) &= \pi^* \\
E(\pi(a^*_{SE}, x)) &= \pi^* \\
E^*(\pi(a^*_{MKT}, x)) &= \pi^*
\end{aligned}$$

The equilibrium in the statistical expectation game involves setting expected inflation, in a statistical sense, equal to target. In contrast, the equilibria in the social welfare and the market-based games involve setting the *market-based* expectation of inflation equal to target. The two kinds of expectations will differ because of the existence of an inflation risk premium:

$$RP(\pi) = E(\pi) - E^*(\pi)$$

This example helps illustrate why policymakers who want to maximize household welfare should favor the use of market-based expectations relative to what are often termed "true" expectations. [10] First, and most important, households are risk averse. As a result, their willingness to substitute losses from one state to another doesn't just depend on the likelihood of those states' occurring - it also depends on the magnitude of losses in that state. If a household expects marginal utility to be high when inflation is high, the household will be willing to pay a lot for insurance against high inflation, even if high inflation is unlikely. If the household expects marginal utility to be low when inflation is low, the household will be willing to pay a lot for insurance against low inflation. Either way, the household's willingness to pay for insurance against particular inflation outcomes – the inflation risk premium – is informative about the welfare associated with those inflation outcomes.Statistical models do not take household resource valuation into account, while market-based probabilities do. Using statistical probabilities to weight resources in different states is, in some sense, equivalent to ignoring discounting when weighting resources in different dates.[11]

Second, households' assessments of the likelihood of various outcomes will typically differ from that of a statistical modeler. These differences may be attributable

[10] The use of the word "true" in this context usually seems to refer to estimates derived from a statistical forecasting model. Note though that different people often have different information and different pre-existing beliefs about the likely future evolution of a given variable of interest. For example, when assessing the odds that inflation will be high or low, different people will often rely on different price changes they have observed or different inflation rates they have experienced during their lives. It is natural for these different people to arrive at different assessments of the probability of various possible future events. There is no clear sense in which one of these assessments is more "true" than any other. In contrast, Ross (2015) shows that there is a unique recovered distribution in a stationary world. Borovička, Hansen, and Scheinkman (2015) extend the analysis of Ross (2015) and establish an additional condition to guarantee that the unique recovered distribution matches the subjective distribution used by investors.

[11] Kitsul and Wright (2013) construct market-based probabilities for inflation. By comparing these probabilities to those from a statistical model, they produce estimates of household resource valuation associated with different outcomes for inflation.

to different information, different beliefs, or the use of unconventional probabilistic modeling. A policymaker who wants to maximize ex-ante social welfare must take these differences into account when answering the basic question.

Some critics of market-based expectations point out that forecasts of the future based on the prices of financial assets don't perform all that well relative to forecasts based on statistical models. This criticism is closely related to the above discussion of market-based probabilities relative to statistically estimated probabilities. The weighting across possible future outcomes embedded in statistical forecasts is an inappropriate benchmark for policymakers because that weighting is based on an inappropriate loss function for policymakers. Generally, statistical forecasts are formed and evaluated using a standard loss function such as mean squared error. But this loss function does not put more weight on a state of the world just because households are more willing to substitute resources toward that state of the world. Hence, this evaluation criterion does not seem particularly relevant for a policymaker who wants to maximize ex-ante social welfare.

5 Conclusions

When policymakers make decisions under uncertainty, they need some way to aggregate the state-contingent net benefits of their decisions across possible outcomes.The question is: how best to do so? This paper provides conditions under which the policymaker maximizes social welfare by using *market-based expectations* to calculate the relevant projections. In contrast, maximizing the statistical expectation of net benefits is generally not equivalent to the maximization of social welfare.

Our main equivalence result assumes that there is a representative agent and that asset markets are complete. It is possible to generalize these conditions in a number of ways. Basically, our result is grounded in an old idea from public economics: prices reflect marginal rates of substitution. In this context, the prices are market-based probabilities and the marginal rates of substitution are with respect to consumption in various states of the world. Of course, this connection between prices and marginal rates of substitution is not literally true in all economic settings.[12] Many authors have argued that the behavior of asset prices is best understood through models that assume that different agents trade in distinct asset markets.[13] Such models sever the exact connection between market-based probabilities and inter-state marginal rates of substitution that lies at the heart of our main result.

Despite these limitations, we see our results as implying that policymakers should view market-based probabilities as a key source of information in their decision-making. For example, consider the problem of an inflation-targeting central bank that we analyzed in Section 4. We agree with Bernanke and Woodford (1997) that it is useful for such central banks to base their decisions on explicit models of inflation. Our message in this paper is that those models should be informed by the behavior of market-based inflation expectations. Without that information, central banks could be led to make policy choices that do not put sufficient weight on households' aversion to inflation risk.

[12]We discuss this issue in considerably more detail in an earlier version of the paper (Feldman, et. al., 2015).

[13]See Guvenen (2009), among others.

References

[1] Armenter, R., 2015. "On the Use of Market-Based Probabilities for Policy Decisions," Federal Reserve Bank of Philadelphia working paper.

[2] Bernanke, Ben S., and Michael Woodford. 1997. "Inflation Forecasts and Monetary Policy." *Journal of Money, Credit and Banking* 29: 653—684.

[3] Borovicka, Jaroslav; Hansen, Lars P., and Jose Scheinkman. 2015. "Misspecified Recovery." NYU working paper.

[4] Cochrane, John. 2011. "Discount Rates," *Journal of Finance* 66: 1047-1108.

[5] Ellison, Martin, and Thomas J. Sargent. 2012. "A Defense of the FOMC." *International Economic Review* 53: 1047—1065.

[6] Feldman, Ron; Heinecke, Ken; Kocherlakota, Narayana; Schulhofer-Wohl, Sam; and Thomas Tallarini. 2015. "Market-Based Probabilities: A Tool For Policymakers," Federal Reserve Bank of Minneapolis working paper.

[7] Guvenen, Fatih. 2009. "A Parsimonious Macroeconomic Model for Asset Pricing." *Econometrica* 77: 1711—1750.

[8] Hetzel, Robert. 1992. "Indexed Bonds as an Aid to Monetary Policy." *Federal Reserve Bank of Richmond Economic Review.*

[9] Kitsul, Yuriy and Jonathan H. Wright. 2013, "The Economics of Options-Implied Inflation Probability Density Functions." *Journal of Financial Economics* 110: 696-711.

[10] Ross, Steve. 2015. "The Recovery Theorem." *The Journal of Finance*, 70: 615–648.

[11] Sumner, Scott. 1995. "The Impact of Futures Targeting on the Precision and Credibility of Monetary Policy." *Journal of Money, Credit and Banking* 27, 89-106.

[12] Woodford, Michael. 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton, NJ: Princeton University Press.