Market-Based Probabilities: A Tool for Policymakers*

January 7, 2015

1 Basics of Policy Choice under Uncertainty

Policymakers make many decisions on behalf of the public. They decide how to spend funds on public works, how to regulate firms and households, and what types of social safety nets to create, among many other examples. Some costs and benefits from these decisions are experienced today. Other costs and benefits are realized in the future. The future is — typically — uncertain, and so the net social benefit of the policymaker's current choice is uncertain.

To understand this point, consider a very large government infrastructure project, such as a dam built to prevent flooding. The costs of the project depend on the costs of inputs such as concrete, which will be purchased in the near term. The benefits of the project are realized in the future and depend on many factors, including the amount of rain in the future. However, policymakers may be uncertain about future weather patterns. The project's net benefits will be higher if future rainfall turns out to be higher, and lower if future rainfall turns out to be lower.

Alternatively, consider a central bank that is deciding how to set monetary policy. The future inflation rate depends on unknown events, such as shocks to oil prices. But the central bank can also influence the future inflation rate by changing the level of monetary accommodation today. The net benefits of providing more monetary accommodation today will be higher if the as-yet-unknown future shocks push inflation below the ideal level, and lower if future shocks push inflation above the ideal level. How should policymakers weigh these different possible outcomes?

As these examples suggest, to make informed choices in the face of uncertainty, policy-

*Federal Reserve Bank of Minneapolis: Ron Feldman, Ken Heinecke, Narayana Kocherlakota, Sam Schulhofer-Wohl, Tom Tallarini.

makers need a systematic way to assess the net benefits of today's decisions across the many possible directions the future may take. This systematic assessment methodology should also consider how households view the range of future costs and benefits. Why? Because policymakers are representing households when they make decisions. Simply put, policymakers act on behalf of households and thus should think about the future as households do.

In the example of building a dam, the policymaker needs a way to weigh the net benefits of the dam when future rainfall and flood risks are not yet known. This assessment will inform the policymaker's decision today. Because the policymaker acts on behalf of households, the assessment should capture households' views about flood risks. Similarly, in the case of monetary policy, the central bank's assessment of the costs and benefits of monetary accommodation should take account of households' views about inflation risks.

We can summarize this discussion in the form of a basic question for policymakers:

In making a current decision, how should the policymaker weigh the net benefits of that decision in one possible future against the net benefits of that decision in another possible future?

The Federal Reserve Bank of Minneapolis answers this question in the following way:

Policymakers should typically weigh the net benefit in a given possible future using what we term the *market-based probability* of that future occurring.[1]

A market-based probability is a weight that financial markets assign to a possible event. These weights are identifiable from asset prices available every day. As investors adjust their expectations, these asset prices change accordingly. As a result, they provide a framework for systematically considering different paths the future may take. It is important that market-

[1] *Market-based probabilities* are also often called *risk-neutral probabilities*.

based probabilities depend on the weights that households attach to different possible futures. This means that policymakers who rely on market-based probabilities will take policy actions that tend to be in line with households' own preferences.

The Minneapolis Fed 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 `http://www.minneapolisfed.org/banking/mpd/`.

The rest of the paper explains our answer in greater detail and increasing complexity for those interested in the technical underpinnings. Section 2 is aimed at readers who have had an exposure to economics and statistics along the lines of undergraduate college classes. The technical analysis in Section 3 is geared toward readers who have postgraduate training in economics or finance.

2 The Economic Intuition for Using Market-Based Probabilities

This section is intended for readers who have had an exposure to economics and statistics along the lines of undergraduate college classes. We will first develop the rationale for informing policymaker decisions with market-based probabilities. We will then respond to common questions about our proposal.

Our basic question asked how policymakers should weigh the net benefits of a decision in one possible future against the net benefits in another possible future. In arriving at our recommendation that policymakers use market-based probabilities, we start from the following premise:

Policymakers should make decisions in the interests of a typical household.

This premise implies that, on the margin, policymakers should weigh the net benefits of a given choice in one possible future against the net benefits in another possible future just as a typical household would. Put another way, a policymaker's marginal willingness to substitute resources in one possible future for resources in another possible future should mimic that of a typical household. We take this view because policymakers are agents of households and should make decisions that, in the household's view, make it better off.

How can we learn about a typical household's willingness to substitute resources in one future for resources in another future? The answer lies in financial market prices. Financial assets differ in terms of what their owner receives in different possible futures. For example, a risk-free U.S. Treasury bond makes the same dollar payments in all possible futures. In contrast, the buyer of a risky corporate bond will receive smaller payments if the future turns out to be one in which the bond issuer faces financial difficulties. Hence, the relative prices of these two financial assets implicitly reveal the willingness of households to substitute resources between two possible futures: a future in which the corporate bond issuer is financially sound (proxied by the risk-free security) and a future in which the issuer is troubled (captured by the higher rates paid by risky securities).

We can readily generalize this basic idea. By using the prices of many financial assets (especially options), we can impute a typical household's marginal willingness to substitute resources in one possible future for resources in a wide variety of other possible futures. When we use the term "market-based probabilities," we are referring exactly to the outcome of this imputation exercise. The policymaker can then make decisions in that typical household's interest by weighting resources in different futures using market-based probabilities.

We now address some commonly expressed concerns about the use of market-based probabilities.

2.1 Market-Based Probabilities Aren't "True" Probabilities

We often receive the comment that policymakers should weight net benefits in a possible future using the "true" probability of that future's occurring. We find the use of the word "true" unclear, but we think it refers to estimates derived from a statistical forecasting model.[2] There are two closely related reasons why policymakers should not use these estimated probabilities to answer the basic policy question.

First, and most important, households' willingness to substitute resources from one possible future to another depends on the relative scarcity of resources in those futures. Thus, a household may be willing to pay a lot for insurance against the possibility of job loss or a flood, even if the household sees the outcome as highly unlikely. Policymakers must take this factor into account when answering the basic question if they hope to act in a way that will improve household well-being. Statistical models do not take household resource valuation into account, whereas market-based probabilities do.[3]

Second, households' assessments of the likelihood of various outcomes may differ from that of a statistical modeler. These differences may be attributable to different information, different beliefs, or the use of unconventional probabilistic modeling. A policymaker who wants to act in the interests of households must take these differences into account when answering the basic question.

[2]Different people often have different information and different preexisting 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. Naturally, these different people will 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 (forthcoming) shows that there is a unique recovered distribution in a stationary world. Borovička, Hansen, and Scheinkman (2014) extend the analysis of Ross (forthcoming) and establish an additional condition to guarantee that the unique recovered distribution matches the subjective distribution used by investors.

[3]Kitsul and Wright (2013) construct market-based probailities 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.

2.2 Market-Based Probabilities Don't Forecast the Future Well

Some critics of market-based probabilities point out that forecasts of the future based on the prices of financial assets do not perform very 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. Thus, our responses are nearly the same.

We reiterate our key point with slightly different words: Policymakers are at least as concerned about the range of possible outcomes a decision might lead to as they are about the realized outcome. So although market-based probabilities may not produce very reliable forecasts, they do consider households' views about which possible future outcomes are particularly bad or particularly good. Put another way, the evaluation criterion of forecasting performance is generally based on 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 act in households' interest.

In a different context, Ellison and Sargent (2012) account for the poor forecasting of the Federal Open Market Committee (FOMC) relative to the Federal Reserve staff's Greenbook forecasts by attributing a concern for robustness on the part of the FOMC. As in the present analysis, Ellison and Sargent rely on the fact that the policymaker is not explicitly trying to produce an optimal forecast, but rather to maximize social welfare.

2.3 There Is No "Typical" Household — Households Are Heterogeneous

In answering the basic question, our premise is that policymakers should act in the interest of a "typical" household. The reality is that households differ in a number of ways. How, then, can we use market-based probabilities to inform policymakers about the views of "typical" households?

The answer lies in the power of financial markets. If households are able to trade a set of assets, then in equilibrium, they are all equally willing, on the margin, to substitute any one of those assets for another asset. As long as the assets differ sufficiently in terms of their payoffs in different possible futures, we can conclude as well that households that participate in financial markets are all equally willing, on the margin, to substitute resources from one possible future for resources in another possible future.

In this sense, even if households are quite different before trading, the act of trading will make them at least marginally identical in terms of how they weight resources in different possible futures. A policymaker who uses market-based probabilities is using that same weighting.

2.4 Diminishment of Private Sector Information Gathering

Market-based probabilities are valuable, in part, because they reflect private sector information about the economy. However, some observers have expressed the concern that policymaker actions based on these market-based probabilities could eliminate the private sector's willingness to gather that information. Roughly speaking, the intuition for this perspective is as follows. Suppose the central bank commits to keep a market-based inflation forecast equal to a fixed number (such as 2 percent). Then the private sector will not be able to make money by out-forecasting the market, and the private sector will have no incentive to gather information about future inflation.

This intuition may be valid in some limited contexts. However, we are considering situations in which a policymaker's action choice has many possible outcomes that are beyond the control of the policymaker. These outcomes influence the prices of options (in the central bank example noted above, options based on the realization of inflation). Regardless of central bank policy, the private sector completely retains its incentive to collect information because that information is valued as a way to beat the market for the options.

2.5 Participation, Nonparticipation and Illiquidity

So far, we have not found the concerns raised about using market-based probabilities compelling. However, observers raise two other concerns that we see as potentially more relevant for policymakers who plan to use market-based probabilities to inform their decisions.

The first is the issue of participation and nonparticipation in asset markets. In particular, we rely heavily on option prices to calculate market-based probabilities. But relatively few people trade in option markets, and some of those who trade are not the U.S. households whose welfare policymakers seek to maximize. How do we know if market-based probabilities, calculated using option prices, reflect the input from U.S. households that do not trade these securities?

The short answer is that we do not. However, we are comforted by two observations. First, households could choose to trade in these markets if they so desired. Indeed, by buying and selling put options and call options, taking either side of a given bet is relatively easy. The nontraders must perceive the gains from undertaking these trades to be small. This reasoning suggests that nontraders' marginal valuations of options relative to other assets are probably not much different from the marginal valuations of people who are trading, even if some of those people are not U.S. households. Second, some option trading occurs at the behest of investors charged with operating as fiduciaries of households, even if the households themselves do not trade.

As we say, these observations are comforting. We do recognize, however, that many households perceive option trading as having nontrivial transactions costs. These perceived costs create the need for some caution in the use of market-based probabilities. Further analysis of the extent to which asset prices reflect U.S. households' valuations would be valuable.

The second issue is that of illiquidity in options markets, which can take two forms. Option trading on some assets is limited in that we often observe meaningful gaps between bids and asks or few trades at out-of-the-money strike prices. As a result, the prices of options may reflect factors beyond investor expectations. (For example, the price could include compensation to investors for holding a security that will be costly to sell.) A policymaker who uses market-based probabilities should take account of these gaps — perhaps by calculating corresponding upper and lower bounds for estimates of market-based probabilities.

Moreover, not much trading at all occurs with options that can inform the tail of the market-based probability distribution. We address this limitation by extrapolating to complete the more extreme part of the distribution. But it is this part of the distribution that is typically of greatest interest to policymakers. Ensuring that extrapolation is robust and subject to review is critical.

3 A Technical Analysis

This section takes a still deeper dive into technical analysis and is intended for readers who are familiar with post-graduate economics or finance. We examine a simple abstract model in which a policymaker makes a decision under uncertainty. Our key result is that, in this setting, maximizing expected net benefits relative to market-based probabilities is equivalent to maximizing expected social welfare relative to households' subjective probabilities. We then consider the extent to which the result generalizes to different models. This latter analysis is closely related to the concerns raised above, but our discussion in this section is more technical.

3.1 Two Policy Games

We consider two policy games that are distinguished by the objective of the policymaker. We begin by describing the common elements of the two games and then describe the different policymaker objectives.

3.1.1 Common Elements

Both games feature households, a policymaker, and 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, that affects outcomes 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. In addition, $B(a, x)$ is measured in units of the consumption good.

One example of a policymaker's problem is that of an inflation-targeting central bank. In that case, we can represent the realization of inflation next period as $\pi = a + x$, where a denotes the level of accommodation provided by the central bank and x represents a shock to inflation such as oil prices or other exogenous events outside the control of the central bank. The central bank has a target for inflation of π^* that is known to all of the agents in the economy. The net benefit function in this case would be $B(a, x) = -(a + x - \pi^*)^2$, which measures the lost consumption associated with the central bank missing its inflation target (see Woodford, 2003).

A somewhat more abstract example is that of a financial regulator. In this case, B again measures net benefits in terms of consumption goods, but now the regulator chooses the size of the capital distributions (dividend payments or share buybacks) that banks may undertake. The postdistribution capital positions of the banking system may affect the likelihood of financial instability in the economy next period. In this framework, the consumption implications of that potential instability are captured in B.

In the trading period, households are identical and their consumption in state n is given by $c(a, x_n) = y(x_n) + B(a, x_n)$, where $y(x)$ is an endogenous endowment that depends on the realization of the state x. Households' expected utility, conditional on policymaker action a, is

$$V(a) = \sum_{n=1}^{N} \pi_n U(y(x_n) + B(a, x_n), x_n). \tag{1}$$

$U(c, x)$ is a possibly state-dependent utility function and is increasing and concave in consumption, and π_n is the households' subjective probability that the state next period will be x_n in the sense of Savage (1954).

3.1.2 Difference between the Games: Policymaker's Objective

The policy games differ with respect to the objective function of the policymaker in the planning period. In the *social welfare* game, the policymaker seeks to maximize social welfare:

$$\sum_{n=1}^{N} \pi_n U(y(x_n) + B(a, x_n), x_n).$$

(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.)

In the *market-based* game, policymakers maximize the market-based expectation of net social benefits. Formally, the policymaker observes the outcome of the trading period. Let q_n denote the implied price today of goods in state n. Now, define

$$q_n^* = \frac{q_n}{\sum_{n=1}^{N} q_n}. \tag{2}$$

Since q_n is the price of goods in state n, $q_n \geq 0$ for all n. As a result, $q_n^* \geq 0$ for all n. In addition, $\sum_{n=1}^{N} q_n^* = 1$. Therefore, $\{q_n^*\}_{n=1}^{N}$ is a probability measure over the states of the world. We will call this the *market-based* probability measure.[4] Given the market-based probability measure, we can define a new expectations operator, E^*, over any random

[4]The $\{q_n^*\}$s are often referred to as *risk-neutral* probabilities, especially in finance.

variable ϕ:

$$E^*[\phi] = \sum_{n=1}^{N} q_n^* \phi_n, \tag{3}$$

where ϕ_n is the realization of the random variable ϕ in state x_n. In the market-based game, the policymaker's objective function is given by

$$E^*[B(a, x)].$$

3.1.3 Equilibrium Equivalence Result

In both of these games, households are identical and so they do not trade in equilibrium. The key aspect of an equilibrium outcome is the policymaker's action choice in the planning period.

In the social welfare game, any equilibrium outcome a^*_{SW} must satisfy the first-order condition

$$\sum_{n=1}^{N} \pi_n MUC_n(a^*_{SW}) \frac{\partial B}{\partial a}(a^*_{SW}, x_n) = 0, \tag{4}$$

where $MUC_n(a) = U_c(y(x_n) + B(a, x_n), x_n)$, the marginal utility of consumption in state n given that the policymaker makes choice a.

In the market-based game, any equilibrium outcome a^*_{MKT} is characterized by the first-order condition

$$\sum_{n=1}^{N} q_n^* \frac{\partial B}{\partial a}(a^*_{MKT}, x_n) = 0. \tag{5}$$

In equilibrium, the households in the trading period know the action choice of the policymaker in the planning period. Furthermore, in equilibrium, households must equate the marginal cost of an extra unit of consumption in each state, q_n, with the marginal benefit of that extra unit of consumption in terms of expected utility, $\pi_n MUC_n(a^*_{MKT})$. Hence:

$$q_n^* = \frac{\pi_n MUC_n(a^*_{MKT})}{\sum_{m=1}^{N} \pi_m MUC_m(a^*_{MKT})}.$$

It follows that any equilibrium outcome in the market-based game is characterized by the first-order condition

$$\sum_{n=1}^{N} \pi_n MUC_n(a^*_{MKT}) \frac{\partial B}{\partial a}(a^*_{MKT}, x_n) = 0 \tag{6}$$

These two first-order conditions (4) and (6) are the same. As long as this equation has a unique solution — which will typically be the case when there is curvature in the utility function or the net benefit function — the equilibrium outcomes in the two games will be the same.

We could have considered a third game in which the policymaker's objective is given by the "true" expectation of net benefits:

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

However, the equilibrium outcome a^* in this game is characterized by the first-order condition

$$0 = \sum_{n=1}^{N} \pi_n \frac{\partial B}{\partial a}(a^*, x_n), \tag{7}$$

which is typically different from the first order conditions (4) and (6). So a^* would *not* be the same as a^*_{SW} or a^*_{MKT}.

To sum up, a policymaker who maximizes social welfare will make the same choices as a policymaker who maximizes the market-based expectation of net benefits. However, the policymaker can solve the social welfare problem only if he knows the net benefit function, $B(a, x)$; the state-dependent marginal utilities, $\{MUC_n(a)\}$; and the household subjective probabilities, $\{\pi_n\}$. The policymaker who seeks to solve for the market-based expectation of next benefits only needs to know the net benefit function B and the state prices q.

Returning to the example of an inflation-targeting central bank in section 3.1, $B(a, x) = -(a + x - \pi^*)^2$, where $a + x = \pi$, the observed rate of inflation. The central bank would choose a so that $E^*[\pi]$ is π^*. In other words, the central bank would set the market-based expectation of inflation (equivalently, the price of a zero coupon inflation swap) equal to the inflation target.

3.2 Possible Concerns

The equivalence between maximizing net expected benefits relative to market-based probabilities and maximizing expected social welfare relative to households' subjective probabilities is based on a particular model. This section discusses several theoretical and practical concerns about the robustness of the equivalence result derived above.

3.2.1 Heterogeneity

The analysis above assumes that all households are identical, both ex ante and ex post. However, policy actions may affect some households differently from others. For example, some households may benefit from a given policymaker's action, whereas some may be made worse off. A common example is the reduction of trade barriers that benefits many households by means of lower prices for the liberalized goods, but results in adverse outcomes for those involved in domestic production of those goods.

Recall that the benefit function $B(a, x_n)$ measures *net* benefits. Therefore, we can assume that any redistributive effects of policy choice a are undone by transfers so that $B^i(a, x_n)$ (where the superscript i indexes the benefit function by agent) has the same sign for all agents in the economy for a given policy-state combination. Taking this idea a step further, we could assume that transfers are such that $B(a, x_n)$ is the same for each agent. In that case, maximizing social welfare would continue to be equivalent to maximizing $E^*[B]$. To see this, generalizing (4) to many agents yields

$$\sum_{i=1}^{I} \omega^i \sum_{n=1}^{N} \pi_n^i MUC_n^i(\hat{a}) \frac{\partial B}{\partial a}(\hat{a}, x_n) = 0, \tag{8}$$

where the ω^i's represent the weights assigned to each agent i by the policymaker and $\widehat{a}$ is the policymaker's optimal choice. As long as they are unconstrained, all agents will set their subjective probability-weighted marginal utilities proportional to the state prices q_n:

$$\frac{q_n}{\xi^i} = \pi_n^i MUC_n^i(\hat{a}), \tag{9}$$

where ξ^i is an agent-specific constant of proportionality. In consequence,

$$\sum_{i=1}^{I} \omega^i \sum_{n=1}^{N} \frac{q_n}{\xi^i} \frac{\partial B}{\partial a}(\hat{a}, x_n) = \sum_{i=1}^{I} \frac{\omega^i}{\xi^i} \left[\sum_{n=1}^{N} q_n \frac{\partial B}{\partial a}(\hat{a}, x_n) \right] = 0. \tag{10}$$

So the equivalence holds.

3.2.2 Limited Participation

Relatively few households trade in the options markets that are used to compute market-based probabilities. As a result, the prices in these markets may not be particularly representative of the population as a whole. In that case, maximizing market-based expected benefits could be equivalent to maximizing the welfare of only the households that participate in options markets – some of which may not be U.S. households. This would be a significant concern if the households that did not participate in options markets were barred from doing so, as in Guvenen (2009). However, to the extent that non-participation is a choice, such a choice would indicate relative valuations of resources state by state that are close to those implied by options markets.

However, in some instances, agents may wish to consume more in a particular state but be unable to do so. In that case, the state price would be less than the scaled, subjective probability-weighted marginal utility for the constrained agent: $q_n < \xi \pi_n MUC_n(\hat{a})$. By maximizing $E^*[B]$, the policymaker would be underweighting the states of the world in which agents are constrained. Conversely, if agents wanted to reduce consumption in some states but could not, the state price would be greater than the scaled, subjective probability-weighted marginal utility for the constrained agent: $q_n > \xi \pi_n MUC_n(\hat{a})$. The policymaker would then overweight these states of the world. Although it would be ideal if policy could alleviate these constraints, they are often the result of asymmetric information or limited commitment and as such may be beyond the scope of the policymaker's ability, especially in the case of monetary policy.[5]

3.2.3 Costly Information Acquisition

Information is costly to obtain. If the information offers no financial return, then agents will not be willing to incur the acquisition costs. Bernanke and Woodford (1997) highlight the potential consequences of this insight for a central bank that chooses policy so as to target private sector inflation forecasts, possibly as reflected in inflation breakevens. If the central bank is viewed as successful in its targeting, the private sector has little incentive to collect relevant information, and the forecasts and prices will not embed socially valuable information. As a result, Bernanke and Woodford argue that a central bank should use a structural economic model rather than private inflation expectations to guide its policy choices.

However, the Bernanke-Woodford analysis does not apply to our setting. The computation of the market-based probability measure intrinsically depends on multiple options with varying strike prices rather than a single inflation forecast. It is true that the policymaker is essentially targeting a particular asset's price to equal zero — namely, the asset with a payoff equal to the marginal net benefit of the policymaker's action. However, the whole basis of the market-based probability approach is that this marginal net benefit asset is not the only one available for trade. Rather, traders can buy and sell a wide variety of assets (options). The policymaker's action does not affect the incentives of traders to acquire information about these other assets. Through arbitrage, this information will get reflected in the price of the marginal net benefit asset itself.

[5] Nonparticipation may be a greater concern with regard to discounting future net benefits. Households that face borrowing constraints also face higher shadow discount rates.

3.2.4 Incompleteness of Asset Markets

If markets are incomplete, there will be multiple market-based probability measures consistent with observed asset prices. In order for the equivalence result to hold, for any action a, the benefit $B(a, x)$ must be spanned by the payoffs of observed assets. This information tells us how much households value each action a. Then, the action that maximizes social welfare is simply the one that has the highest value for households. For example, consider the case in which $N = 3$, but there are only two assets.[6] The first asset, z_1, pays one unit of consumption when the state is x_1 and zero in the other two states. The second asset, z_2, pays one unit of consumption when the state is x_2 or x_3 and zero in the first state. As long as the benefit function satisfies $B(a, x_2) = B(a, x_3)$ for all choices[7] of a, then it is spanned by the two assets and the policymaker can determine the optimal action by maximizing $E^*[B]$.

Even without spanning, a weaker version of the equivalence result may be obtained in the form of a necessary condition for social optimality. As noted above, with incomplete markets, there is a set of market-based probability measures consistent with observed asset prices. This means that for any policymaker action, there is a set S of market-based expectations of the marginal net benefits associated with that action. We can show that, if it is common knowledge that the policymaker will choose the action that maximizes social welfare, then zero lies in the set S. Put another way, if no element of S is zero, then we know that households believe that the policymaker is not going to make a socially optimal choice. Intuitively, one element of S is defined by the marginal utilities of consumption of the typical household, and that element generates the first-order conditions of the social welfare maximization problem.

[6] In particular, $y(x_n)$ has three distinct values.

[7] This incomplete markets result is based on the assumption that all households are identical.

3.2.5 Policy Uncertainty

Up to now, we have assumed that the policymaker and households are symmetrically informed. This assumption implies that, during the trading period, households know what the policymaker will choose in the planning period. In other words, households face no policy uncertainty. But what if the policymaker has information available in the planning period that the households do not know in the trading period? In this subsection, we sketch an argument that the main equivalence result will generalize, as long as asset markets are sufficiently complete.

More specifically, suppose that the households' endowment, the net benefit function, and the utility function depend not only on a and x but also on a second random variable, Z. Suppose too that Z is realized after households trade contingent claims but before the policymaker acts. Because U, y, or B depends on the realization z of Z, the policymaker's optimal action for any given x will depend on z. Thus, in the trading period, the households cannot be certain how the policymaker will act for any given realization of x.

This uncertainty does not affect our results if markets are complete. Under complete markets, households trade claims in the trading period that are contingent on the realization of both x and z. Again, consider two games distinguished by the policymaker's objective. In the social welfare game, for each realization z of Z, the policymaker's strategy $a^*_{SW}(z)$ solves the problem

$$\max_a \sum_{n=1}^{N} \pi_n(z) U(y(x_n, z) + B(a, x_n, z), x_n, z),$$

where $\pi_n(z)$ is the subjective probability of x_n, conditional on Z equaling z. In the market-based game, at the time of trading, the households know the policymaker's strategy $a(z)$ as a function of the (as yet unknown) random variable z. Hence, the market-based probability $q^*_n(z)$ of x_n, conditional on Z equaling z, is given by

$$q^*_n(z) = \sum_{n=1}^{N} \pi_n MUC_n(a^*_{MKT}(z); z), \tag{11}$$

where $MUC_n(a; z) = U_c(y(x_n, z) + B(a, x_n, z), x_n, z)$. For any z, the equilibrium strategy a^* is a solution to the following problem:

$$\max_a \sum_{n=1}^{N} q_n^*(z) B(a, x_n, z).$$

Given these assumptions, a version of (4) still describes the policymaker's equilibrium strategy a^*_{SW} in the social welfare game. Similarly, a version of (6) describes the policymaker's equilibrium strategy a^*_{MKT} in the market-based game. By using these first-order conditions, we can obtain the same equivalence result: the Nash equilibrium outcome in the social welfare game is the same as in the market-based game.

As we noted above, we can obtain a weaker version of our equivalence result even if markets are incomplete with respect to z. For example, suppose that households can trade a complete set of claims contingent on x, but no claims contingent on z. Then, there is a set of market-based probability measures that are consistent with the prices of the traded claims. We can show that, if the policymaker uses the socially optimal strategy a^*_{SW}, then there is some market-based probability measure that implies that the market-based expectation of the marginal net benefit of a is equal to zero. Again, the intuition is that some market-based probability measure is associated with the marginal utility of consumption of the houeholds.

This last result — the socially optimal strategy maximizes net benefits relative to some set of market-based probabilities — is more informative when markets are more complete and there is less policy uncertainty. In the limiting cases of no policy uncertainty (Z has only one possible realization) or complete markets (claims are contingent on both x_n and z), only one market-based probability measure is consistent with the observed prices, and hence maximizing net benefits relative to that measure is not only necessary but also sufficient for optimality. At the other extreme, if there are no contingent claims on either x or z, any set of nonnegative numbers that sums to one is consistent with the observed prices, and no restriction comes from limiting ourselves to market-based probabilities. Thus, our claim that policymakers should set policy so as to maximize expected benefits relative to market-based probabilities is more useful when there is less policy uncertainty.[8]

3.3 Conclusion

Policy decisions affect the economy with a lag. Hence, policymakers need some way to gauge the relative likelihoods of future events. For an inflation-targeting central bank, the likelihoods of deflation on the one hand or high inflation on the other are important inputs to the policy-setting process. For a financial regulator, the likelihood of significant financial instability is needed to assess the risks associated with bank capital distributions. However, our analysis suggests that the typical approach of trying to discern the "true" probability of events is typically inappropriate. Instead, policymakers should base their decisions on market-based probabilities or their equivalents. These probabilities encode households' ex ante marginal valuations of resources in different states of the world as well as their subjective likelihood of those states. An increase in the market-based probability of an outcome such as deflation could indicate that households consider it as more likely, or it could indicate that the costs associated with deflation have risen. Both of these changes should matter for a policymaker. And they do if the policymaker maximizes expected benefits relative to those market-based probabilities.

References

Bernanke, Ben S., and Michael Woodford. 1997. "Inflation Forecasts and Monetary Policy." *Journal of Money, Credit and Banking* 29(4): 653–684.

Borovička, Jarolslav, Lars Peter Hansen, and José A. Scheinkman. 2014. "Misspecified Recovery." manuscript.

[8]From a normative perspective, our simple model certainly implies that policymakers should credibly reveal z if that can be done.

Ellison, Martin, and Thomas J. Sargent. 2012. "A Defense of the FOMC." *International Economic Review* 53(4): 1047–1065.

Guvenen, Fatih. 2009. "A Parsimonious Macroeconomic Model for Asset Pricing." *Econometrica* 77(6): 1711–1750

Kitsul, Yuriy and Jonathan H. Wright. 2013, "The Economics of Options-implied Inflation Probability Density Functions." *Journal of Financial Economics* 110(3): 696-711.

Ross, Steve. Forthcoming. "The Recovery Theorem." *Journal of Finance.*

Savage, Leonard J. 1954. *The Foundations of Statistics.* New York: Wiley.

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