

Polarization and Ambiguity[†]

By SANDEEP BALIGA, ERAN HANANY, AND PETER KLIBANOFF*

A number of voters are in a television studio before a US presidential debate. They are asked the likelihood that the Democratic candidate will cut the budget deficit, as he claims. Some think it is likely and others unlikely. The voters are asked the same question again after the debate. They become even more convinced that their initial inclination is correct. A similar phenomenon can arise in financial markets. “Bulls” and “bears” have different beliefs. On seeing the same macroeconomic forecasts, they become more bullish and bearish, respectively. Individuals observe the same evidence, and yet their beliefs move in opposite directions and end up further apart.

Similar polarization of beliefs has been documented in experiments. For example, Darley and Gross (1983) randomize subjects into different groups. They show one group evidence suggesting a child is from a high socioeconomic background; another that she is from a low socioeconomic background. The former predict the child’s reading abilities are higher than the latter. The groups then watch a film of the child taking an oral test on which she answers some questions correctly and others incorrectly. Those who received the information that the child came from a high socioeconomic background, rate her abilities higher than before; those who received the information indicating she came from a low socioeconomic background rate her lower than before. Thus, the common evidence—the film—leads beliefs to polarize.

We follow Dixit and Weibull (2007) in defining polarization as follows: suppose that, prior to observing a common signal, two individuals have different beliefs and individual A ’s belief first-order stochastically dominates individual B ’s. Their beliefs exhibit *polarization* if, after observing the common signal, individual A ’s posterior dominates his prior and individual B ’s prior dominates his posterior.

Consider two individuals who agree on the probability of each signal conditional on an underlying parameter and use Bayes’ rule to update their beliefs. We show polarization cannot occur. As individuals share the same theory connecting parameters to signals, a given signal increases one individual’s belief if and only if it increases the other’s (Theorem 1).

*Baliga: Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208 (e-mail: baliga@kellogg.northwestern.edu); Hanany: Department of Industrial Engineering, Tel Aviv University, Tel Aviv 69978, Israel (e-mail: hananye@post.tau.ac.il); Klibanoff: Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208 (e-mail: peterk@kellogg.northwestern.edu). Some previous versions of this paper had the title “Ambiguity, Dynamic Consistency and Behavioral Phenomena.” We thank Subir Bose, Jeff Ely, Itzhak Gilboa, Anna Gumen, Christian Kellner, Sujoy Mukerji, Marco Ottaviani, Emre Ozdenoren, Wolfgang Pesendorfer, Ben Polak, Jean-Marc Tallon, three anonymous referees, and a number of seminar and conference audiences for helpful comments and discussion. We especially thank Eran Shmaya for his help with the proof of Theorem 1. Daniel Garrett and Luciano Pomatto provided excellent research assistance. This research was partially supported by Grant No. 2006264 to Hanany and Klibanoff from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel. The authors have no financial or other material interests related to this research to disclose.

[†]Go to <http://dx.doi.org/10.1257/aer.103.7.3071> to visit the article page for additional materials and author disclosure statement(s).

In contrast, we show that polarization can occur as an optimal response to ambiguity aversion (i.e., aversion to subjective uncertainty about probabilities). We do this within a simple prediction model: an individual must predict the value of a parameter that determines the distribution of a random variable. The individual views the parameter as ambiguous and is ambiguity averse. He observes a signal that can inform his predictions. His payoff is decreasing in the squared difference between his prediction and the parameter. This is a standard model apart from ambiguity aversion.

How does ambiguity aversion affect behavior? An individual is exposed to ambiguity when the expected payoff to his strategy varies with probabilities over which he is uncertain. Different strategies may involve different exposure to ambiguity. Suppose there are just two possible parameter values, 0 and 1. If the individual predicts $\frac{1}{2}$, then the difference between the prediction and the parameter is the same no matter what the parameter value. Thus, this prediction strategy completely hedges against ambiguity, i.e., removes any exposure to the ambiguity about the parameter. However, if the individual predicts 1, the squared difference is much higher when the parameter is 0 than when it is 1, exposing the individual to ambiguity. An ambiguity-averse individual will tend to favor strategies that reduce exposure. Reducing exposure is not the only concern of such an individual—for example, the more weight his beliefs place on the parameter equaling 1, the higher his optimal prediction.

Simple strategies such as “predict $\frac{1}{2}$ ” or “predict 1” are generally not optimal when a signal is forthcoming, as these strategies fail to condition on anticipated information. An optimal contingent strategy will make the prediction an increasing function of the observed signal likelihood ratio. We focus on individuals who form an ex ante optimal contingent strategy (i.e., optimal assuming full commitment to that strategy once chosen) and who are indeed willing to carry it out after each possible contingency. Such an individual is said to be *dynamically consistent*.¹ Dynamic consistency is assumed in almost all standard economic analysis. Any theory of updating that is not dynamically consistent will lead to worse outcomes as evaluated by ex ante welfare. Also, studying the dynamically consistent case identifies the key effects leading to polarization that apply even when substantial dynamic inconsistency may be present (see Section IIIB).

Suppose the individual concludes that “predict $\frac{3}{4}$ if the signal is high, $\frac{1}{2}$ if the signal is medium, and $\frac{1}{4}$ if the signal is low” is his optimal strategy. This strategy leaves him partially exposed to the ambiguity about the parameter. Notice, however, that this exposure varies with the signal received. He is not exposed to ambiguity after a medium signal but is exposed after a high or low signal. Under ambiguity aversion, the greater exposure to ambiguity after seeing a high or low signal may lead to an increased desire to hedge against this ambiguity, while the lack of exposure after seeing a medium signal may diminish the value of hedging. These changed hedging motives, ceteris paribus, could lead the individual to want to depart from the ex ante optimal strategy. We call this the *hedging effect*. There is also a more standard effect having nothing to do with ambiguity attitude. After a signal is realized, the likelihoods of this signal are no longer relevant for

¹See, e.g., Hanany and Klibanoff (2007, 2009) for such an approach to modeling ambiguity-averse individuals and for discussion and references to alternative approaches.

optimality going forward. We call this the *likelihood effect*. Because of this effect, after seeing the signal, if beliefs were not updated to incorporate that signal's likelihoods, the individual might want to depart from the ex ante optimal strategy. Dynamically consistent updating must neutralize both the hedging and the likelihood effects of the signal on the incentives of an ambiguity-averse individual. Bayesian updating counterbalances only the likelihood effect. The presence of the hedging effect leads dynamically consistent updating to necessarily depart from Bayes' rule under ambiguity aversion.² Importantly, the hedging effect may alter the direction of updating. Moreover, the hedging effect (but not the likelihood effect) depends on the ex ante optimal strategy which, in turn, is influenced by the individual's beliefs before observing the signal. Through this chain of reasoning, beliefs can influence the direction of updating. This allows for the possibility of polarization. We study when polarization does and does not occur.

Even under ambiguity aversion, we show that polarization does not occur after observing the highest or the lowest signals (Theorem 2). Thus, polarization is a possibility only at signals with an intermediate likelihood ratio. We can offer a particularly clean result if the intermediate signal is neutral (i.e., has equal probability under both parameter values). Then the hedging effect is the only reason to update beliefs—there is no likelihood effect. We show that individuals with sufficiently extreme and opposite beliefs display polarization after observing a common neutral intermediate signal (Theorem 3). When there are exactly three possible signals and there is constant relative ambiguity aversion, we provide necessary and sufficient conditions for polarization (Theorem 4) that apply even when the intermediate signal is not neutral.

All of the above results apply whether or not the two individuals have the same degree of ambiguity aversion as long as initial beliefs differ. Finally, even if individuals have the same beliefs, as long as several conditionally independent observations are available then if they observe different private signals before observing a common signal, they can have different beliefs by the time they see the common signal and, at that point, our above results apply. We now turn to our model. Related literature is discussed at the end of the article.

I. The Model and Benchmark Result

Consider an individual who is concerned with the value of a parameter $\theta \in \Theta \subset \mathbb{R}$. His beliefs are given by a full-support prior μ . To help inform the individual about θ , conditionally independent observations from a random variable X given θ may be available. This random variable has distribution π_θ and takes values in a finite set \mathcal{X} such that each $x \in \mathcal{X}$ has $\pi_\theta(x) > 0$ for some $\theta \in \Theta$. For example, θ might indicate a child's reading ability, while π_θ might be the distribution of scores on a reading test for a child with that ability.

²In contrast, for an expected utility maximizer, dynamic consistency requires that subjective beliefs are updated using Bayes' rule, thus ruling out polarization.

We assume that Θ is finite and, without loss of generality, index Θ so that $\theta_1 < \theta_2 < \dots < \theta_{|\Theta|}$. A distribution $\hat{\eta}$ (first-order) stochastically dominates $\check{\eta}$ if

$$\sum_{i=1}^k \check{\eta}(\theta_i) \geq \sum_{i=1}^k \hat{\eta}(\theta_i) \text{ for all } k \in \{1, 2, \dots, |\Theta|\}.$$

The dominance is strict if at least one of these inequalities is strict. We adopt the following definition of polarization due to Dixit and Weibull (2007):

DEFINITION 1: Fix two individuals with beliefs $\check{\eta}$ and $\hat{\eta}$ over Θ and with common support such that $\hat{\eta}$ stochastically dominates $\check{\eta}$. After they both observe a signal $x \in \mathcal{X}$ whose likelihood given $\theta \in \Theta$ is $\pi_\theta(x)$, we say that polarization occurs if and only if the resulting posterior beliefs, $\check{\nu}$ and $\hat{\nu}$, respectively, lie further apart, i.e., $\check{\eta}$ stochastically dominates $\check{\nu}$ and $\hat{\nu}$ stochastically dominates $\hat{\eta}$ with at least one dominance strict.

Starting from two beliefs, one higher than the other, polarization occurs when the observation of a common signal leads the higher belief to move higher and the lower belief to move lower. The requirement that one initial belief be higher than the other is essential to rule out the possibility that posteriors move in opposite directions but toward each other. If initial beliefs were not ranked, one were updated upward and the other downward, it would not be clear whether beliefs had moved toward or away from one another. Even if the requirement that initial beliefs be ranked were somehow relaxed, the results of this article would continue to hold. Our benchmark impossibility result (Theorem 1) does not rely on the initial beliefs being ranked. Additionally, in the setting we use in the next section to show polarization under ambiguity aversion, beliefs can always be ranked by dominance. In more complex settings, this aspect of defining polarization may matter. We favor Definition 1 because whenever it identifies polarization, there is no doubt that beliefs are moving away from one another.

The following result shows that, in this setting, under Bayesian updating, irrespective of any nonbelief aspect of preference, polarization cannot occur. The theorem and proof formalize the intuitive statement that, as long as their priors share the same support, if two individuals who use Bayes' rule see the same observation and agree on the probability of each observation conditional on the parameter, it is impossible for them to update in opposite directions in the sense of first-order stochastic dominance. All proofs are in the online Appendix.

THEOREM 1: Fix two individuals as in Definition 1. Polarization cannot occur if they use Bayesian updating.

II. Polarization and Ambiguity

For the remainder of the article, suppose the individual's goal is to predict the value of the parameter θ . For tractability, we assume $\Theta = \{0, 1\}$ so there are just two possible parameter values. Accordingly, throughout this section let μ denote

$\mu(\theta = 1)$. For example, $\theta = 1$ might indicate a child has high reading ability and the prediction $\alpha \in [0, 1]$ might be interpreted as a probability that the child has high reading ability. We make the standard assumption that the payoff to a prediction α is given by quadratic loss; i.e., $-(\alpha - \theta)^2$. To avoid tedious corner cases, we assume π_θ has full support for each θ . The individual can condition his prediction on an observation $x \in \mathcal{X}$ drawn from the random variable X .³ Hence, his strategy is a function $\alpha : \mathcal{X} \rightarrow \mathbb{R}$. We use uppercase X and lowercase x to distinguish between the random variable and its realization, respectively.⁴ We assume the individual views θ as ambiguous, is risk neutral and evaluates prediction strategies according to ambiguity averse smooth ambiguity preferences (Klibanoff, Marinacci, and Mukerji 2005). Specifically, any prediction strategy is evaluated according to the concave objective function

$$E_{(\mu, 1-\mu)}\phi[E_{\pi_\theta}(-(\alpha(X) - \theta)^2)],$$

where ϕ is increasing, concave, and continuously differentiable, E is the expectation operator (with respect to the subscripted distribution), $(\mu, 1 - \mu)$ is the belief about θ and π_θ is the distribution of X given θ . Observe that if ϕ is linear (i.e., ambiguity neutrality), the objective function reduces to expected quadratic loss. The concavity of ϕ reflects ambiguity aversion. We will sometimes additionally assume constant relative ambiguity aversion $\gamma \geq 0$, in which case $\phi(u) = \frac{-(-u)^{1+\gamma}}{1+\gamma}$ for $u \leq 0$.

The optimal strategy $\alpha^*(x)$ is the unique solution to the first-order conditions:

$$(1) \quad \frac{\alpha^*(x)}{1 - \alpha^*(x)} \frac{\phi'[E_{\pi_0}(-(\alpha^*(X))^2)]}{\phi'[E_{\pi_1}(-(1 - \alpha^*(X))^2)]} \frac{\pi_0(x)}{\pi_1(x)} = \frac{\mu}{1 - \mu}$$

for each $x \in \mathcal{X}$.

The difference from the usual prediction problem with ambiguity neutrality is the presence of the term $\frac{\phi'[E_{\pi_0}(-(\alpha^*(X))^2)]}{\phi'[E_{\pi_1}(-(1 - \alpha^*(X))^2)]}$ on the left-hand side of (1). Under ambiguity aversion, ϕ is concave, and this term reflects the desire to hedge or reduce the variation in expected payoffs as a function of the ambiguous parameter θ . Ambiguity aversion ensures that when expected payoffs across the θ s differ, the ϕ' ratio pushes the optimal prediction strategy in the direction of equalizing them by moving the predictions toward the θ with the lower expected payoff. This is the manifestation of the value that ambiguity-averse individuals place on hedging against ambiguity. For this reason, we call this ϕ' ratio the *hedging motive*. It compares the marginal value of an extra unit of expected utility when $\theta = 0$ to the marginal value when $\theta = 1$. When these expected payoffs are equal (i.e., a perfect hedge) the hedging motive equals one. Values above (below) one reflect a stronger (weaker) desire

³Extensions of our results to the case of multiple conditionally independent observations may be found in the working paper version (Baliga, Hanany, and Klibanoff 2013).

⁴For example, $\alpha(X)$ is a random (since the observation is yet to be realized) prediction, while $\alpha(x)$ is the realized prediction.

to shift expected payoff from $\theta = 1$ to $\theta = 0$, i.e., to hedge by adjusting the prediction strategy α^* downward.

We have the following useful implication of (1) that is true independent of ambiguity attitude: for any $x, y \in \mathcal{X}$,

$$(2) \quad \frac{\alpha^*(x)}{1 - \alpha^*(x)} \frac{\pi_0(x)}{\pi_1(x)} = \frac{\alpha^*(y)}{1 - \alpha^*(y)} \frac{\pi_0(y)}{\pi_1(y)}.$$

The intuition for this equality is the standard one concerning equating marginal rates of substitution across signal realizations.

In a prediction problem, *updating* maps beliefs about θ and new observations to posterior beliefs about θ . *Dynamically consistent updating* is updating that preserves the optimality of the contingent strategy $\alpha^*(x)$ after the observation is realized (i.e., *ex ante* optimal updating).⁵ Let ν_x denote the posterior probability of $\theta = 1$ after observing x . Dynamically consistent updating is equivalent to these posteriors ν_x satisfying

$$(3) \quad \frac{\alpha^*(x)}{1 - \alpha^*(x)} \frac{\phi'[-(\alpha^*(x))^2]}{\phi'[-(1 - \alpha^*(x))^2]} = \frac{\nu_x}{1 - \nu_x}$$

for all $x \in \mathcal{X}$. Note that (3) is simply the first-order condition of the continuation prediction problem, evaluated at the *ex ante* optimal strategy $\alpha^*(x)$, after x has been realized and assuming beliefs at that point are ν_x . It therefore guarantees that $\alpha^*(x)$ remains optimal upon seeing the observation. After the next result, we describe the difference, under ambiguity aversion, between dynamically consistent updating and Bayesian updating and show how the former allows polarization. First we show that several natural properties that hold under ambiguity neutrality continue to hold under ambiguity aversion:

PROPOSITION 1: (i) *For each realization $x \in \mathcal{X}$, the optimal prediction $\alpha^*(x)$ is an increasing function of μ (the prior probability of $\theta = 1$) and of the likelihood ratio $\frac{\pi_1(x)}{\pi_0(x)}$. The best constant (i.e., not varying with x) prediction is also an increasing function of μ ;*

(ii) *The optimal prediction after observing x is an increasing function of the posterior ν_x . Furthermore, $\nu_x \geq \mu$ if and only if the optimal prediction after observing x and the (ex ante) best constant prediction are similarly ordered. Under dynamically consistent updating, the optimal prediction after observing x is $\alpha^*(x)$.*

Proposition 1 implies that polarization as defined in terms of beliefs is equivalent to polarization in actions (here, predictions). A common signal moves optimal actions further apart and in opposite directions exactly when that signal moves beliefs further apart and in opposite directions. To see this, observe from the first

⁵For a thorough discussion and analysis of dynamically consistent updating under ambiguity aversion see Hanany and Klibanoff (2009).

part of the proposition that prior beliefs have the same order as the respective best predictions made without using the observation. From the second part of the proposition, posterior beliefs after a common signal compare to the prior beliefs in the same way as the optimal predictions after a common signal compare to the best predictions made without using the signal. Combining these yields the equivalence.

For the remainder of this section we assume dynamically consistent updating. Suppose a signal x is observed. Dynamic consistency requires that the optimal prediction strategy after this observation also be the optimal prediction strategy ex ante contingent on observing x . As emphasized above, under ambiguity aversion, the optimal prediction strategy is partly driven by the desire to hedge. Before the signal is realized, the hedging motive is as in (1). However, after observing x , the interim hedging motive is as in (3). These hedging motives are typically *not* equal. As we highlighted in the opening section, the individual’s hedging motive changes since he no longer needs to account for variation in his expected payoffs induced by the signal realization. To carry out the optimal prediction strategy, dynamically consistent updating departs from Bayesian updating in a way that exactly offsets this hedging effect. We use this to offer a characterization of the direction of dynamically consistent updating:

PROPOSITION 2: *The posterior ν_x is above/equal to/below the prior μ if and only if*

$$(4) \quad \frac{\phi'[-(\alpha^*(x))^2] \pi_1(x)}{\phi'[-(1 - \alpha^*(x))^2] \pi_0(x)} \geq \frac{\phi'[E_{\pi_0}(-(\alpha^*(X))^2)]}{\phi'[E_{\pi_1}(-(\alpha^*(X))^2)]}.$$

Notice that in addition to the hedging motive terms, (4) includes the likelihood ratio for the newly observed signal. This term reflects the likelihood effect referred to in the opening section. Thus (4) formalizes the statement that dynamically consistent updating offsets both the hedging and likelihood effects (while Bayesian updating offsets only the likelihood effect). The condition in (4) is not always easy to apply, as it involves the endogenously determined optimal strategy α^* . Nevertheless, we can prove some general properties of updating directly from this inequality:

Observing a signal x^H with the highest likelihood ratio (i.e., $x^H \in \arg \max_{x \in \mathcal{X}} \frac{\pi_1(x)}{\pi_0(x)}$) always leads to updating upwards and observing a signal, x^L , with the lowest likelihood ratio always leads to updating downwards. One implication is that polarization cannot occur after “extreme” common signals.

THEOREM 2: *After observing x^H (x^L), the posterior ν_x is above (below) the prior μ . It is also above (below) the Bayesian update of μ given x^H (x^L).*

Remark 1: *If the signal is informative, so that $\frac{\pi_1(x^H)}{\pi_0(x^H)} > 1 > \frac{\pi_1(x^L)}{\pi_0(x^L)}$, then above (below) in the statement of the theorem may be replaced by strictly above (strictly below).*

We turn to our main positive results. Suppose there are two individuals with beliefs $\hat{\eta} > \check{\eta}$ (we continue our abuse of notation and denote $\hat{\eta}(\theta = 1)$ and $\check{\eta}(\theta = 1)$ by $\hat{\eta}$ and $\check{\eta}$, respectively). In all other respects, the individuals are equivalent. If they are ambiguity neutral, we know that polarization is impossible, from Theorem 1. If they both observe the same extreme signal, we know they will update in the same direction and thus fail to polarize, from Theorem 2. So, assume the individuals are ambiguity averse and there are at least three signals with distinct likelihood ratios. Thus, there is at least one intermediate (i.e., nonextreme) signal, and these are the only signals after which individuals' beliefs can possibly exhibit polarization. When and why can polarization occur?

Suppose an individual observes a signal, x^M , with intermediate likelihood ratio. Substituting for predictions $\alpha^*(x)$, $x \neq x^M$ using (2), inequality (4) becomes

$$(5) \quad \frac{\phi' [-(\alpha^*(x^M))^2]}{\phi' [-(1 - \alpha^*(x^M))^2]} \frac{\pi_1(x^M)}{\pi_0(x^M)}$$

$$\lesseqgtr \frac{\phi' \left[-(\alpha^*(x^M))^2 \sum_{y \in \mathcal{X}} \pi_0(y) \left(\frac{\pi_1(y)}{\pi_0(y)} \right) \right]^2}{\phi' \left[-(1 - \alpha^*(x^M))^2 \sum_{y \in \mathcal{X}} \pi_1(y) \left(\frac{\pi_1(x^M)}{\pi_0(x^M)} \right) \right]^2} \cdot$$

The direction of this inequality determines the direction of updating. The connection between $\alpha^*(x^M)$ (and thus beliefs, since $\alpha^*(x^M)$ is increasing in beliefs by Proposition 1), and the direction of this inequality may be quite complex. It is simplest in the case where the signal is not only intermediate but also *neutral* (i.e., $\pi_0(x) = \pi_1(x)$). In the theorem below, we show that when $\alpha^*(x^M)$ (and thus belief) for one individual is close to 0 and for another is close to 1, polarization occurs after they commonly observe a neutral signal.

THEOREM 3 (Polarization and Ambiguity): *Assume there is a neutral signal, at least one informative signal, and twice continuously differentiable ϕ with $\phi'' < 0 < \phi'$. Polarization occurs after commonly observing a neutral signal if belief $\hat{\eta}$ is sufficiently close to 1 and belief $\check{\eta}$ is sufficiently close to 0.*

Sketch of Proof (for the full proof see the online Appendix): When $\alpha^*(x^M)$ is close to 0, if $\theta = 0$ then predictions will be close to perfect, both interim and ex ante. Since payoffs are relatively insensitive to small changes in predictions in the neighborhood of perfection, any differences in the interim and ex ante expected payoffs when $\theta = 0$ (i.e., any differences in the arguments of ϕ' in the numerators on each side of (5)) will be very small and will have minimal influence on updating

(since $-\frac{\phi''}{\phi'}$ is finite).⁶ In contrast, if $\theta = 1$, predictions close to 0 will be very costly and small improvements in those predictions would be valuable. Therefore, (since $-\frac{\phi''}{\phi'}$ is nonzero) it is the differences in interim and ex ante expected payoffs when $\theta = 1$ that drive the comparison of hedging motives when predictions are close to 0. Differentiating the arguments of the ϕ' terms in the denominators with respect to $\alpha^*(x^M)$, and evaluating at $\alpha^*(x^M) = 0$, yields that the ex ante expected payoff when $\theta = 1$ is higher than the interim payoff when $\theta = 1$ if and only if the expected likelihood, $\sum_{y \in X} \pi_1(y) \frac{\pi_1(y)}{\pi_0(y)}$, is higher than the realized likelihood, $\frac{\pi_1(x^M)}{\pi_0(x^M)}$. This comparison reflects the fact that the predictions $\alpha^*(y)$ optimally move toward 1 by an amount proportional to the likelihood $\frac{\pi_1(y)}{\pi_0(y)}$ so the expected or realized likelihoods reflect the expected or realized improvements in the prediction when $\theta = 1$. Notice that this expected likelihood is always larger than 1 because of the complementarity between the π_1 terms, so that if x^M is a neutral signal this condition will be satisfied.

As a result, when the signal likelihood $\frac{\pi_1(x^M)}{\pi_0(x^M)}$ is below $\sum_{y \in X} \pi_1(y) \frac{\pi_1(y)}{\pi_0(y)}$, for all sufficiently low beliefs η (so that $\alpha^*(x^M)$ is sufficiently close to 0) the hedging motive is bigger ex ante than after seeing the signal, and so updating will be shaded downward compared to Bayesian updating. Similar reasoning for $\alpha^*(x^M)$ close to 1 shows that when $\frac{\pi_1(x^M)}{\pi_0(x^M)}$ lies above $\frac{1}{\sum_{y \in X} \pi_0(y) \frac{\pi_0(y)}{\pi_1(y)}}$, for sufficiently high η , updating will be shaded upward compared to Bayesian updating. When the signal is neutral, Bayesian updating is flat, so these arguments imply updating will be downward when belief is sufficiently low and upward when belief is sufficiently high, generating polarization.

Remark 2: In this argument, the important aspect of quadratic loss is that the marginal payoff to improving a prediction is diminishing in the quality (i.e., closeness to the truth) of the prediction and vanishes at perfection. Any payoff function of the form $\psi(|\alpha - \theta|)$ where $\psi : [0, 1] \rightarrow R$ is a twice continuously differentiable function satisfying $\psi'(0) = 0$ and $\psi'' < 0$ will yield a similar result.

A. Threshold Rules for Updating

To further investigate when polarization occurs, we turn to a particularly clean structure for determining the direction of updating. Given an observation x , say that updating follows a *threshold rule* if there is a threshold $\tau \in [0, 1]$ such that all beliefs above the threshold are updated upward and those below the threshold are updated downward. Under ambiguity neutrality, the threshold is always degenerate—a given observation x either leads all priors to be updated upward or all priors to be updated downward depending on how the likelihood ratio $\frac{\pi_1(x)}{\pi_0(x)}$ compares to 1. In contrast,

⁶The role of assuming $\phi'' < 0 < \phi'$ is to ensure that both the hedging motive, $\frac{\phi'[-\alpha^2]}{\phi'[-(1-\alpha)^2]}$, and ambiguity aversion (as measured by $-\frac{\phi''}{\phi'}$, the coefficient of (absolute) ambiguity aversion (see Klibanoff, Marinacci, and Mukerji 2005)), are bounded away from zero and infinity.

under ambiguity aversion, updating may follow a nontrivial threshold rule. In the online Appendix, we provide a characterization of when updating follows a threshold rule (Proposition 5). To provide an explicit description, for the remainder of this section, we specialize by assuming constant relative ambiguity aversion and that there are exactly three distinct likelihood ratios associated with signals. Under these conditions, we show that updating always follows a threshold rule and we can explicitly derive the thresholds. By Theorem 2, all beliefs are updated in the same direction after extreme signals. Hence, we study thresholds given the intermediate signal.

Polarization is obviously impossible if two individuals have the same beliefs and have the same degree of ambiguity aversion. If, however, there is heterogeneity on either dimension, individuals may exhibit polarization when they observe a common signal. Theorem 4 and Proposition 3 characterize the conditions for a signal to lead to polarization when there is heterogeneity across individuals in beliefs and/or ambiguity aversion.

THEOREM 4: *Assume constant relative ambiguity aversion and exactly three distinct likelihood ratios. There exist $\hat{\tau}, \check{\tau} \in [0, 1]$ such that polarization occurs after commonly observing a signal with the nonextreme likelihood ratio if and only if belief $\hat{\eta} \geq \hat{\tau}$ and belief $\check{\eta} \leq \check{\tau}$ with at least one inequality strict.*

Notice that whenever the thresholds satisfy $\hat{\tau} < 1$ or $\check{\tau} > 0$, there exist beliefs that generate polarization. The theorem relies on the following proposition establishing that updating follows a threshold rule. The proposition is proved by explicitly constructing the threshold.

PROPOSITION 3: *Assume constant relative ambiguity aversion $\gamma > 0$ and exactly three distinct likelihood ratios. The posterior ν_{x^m} is above/equal to/below the prior μ when the latter is above/equal to/below a threshold $\tau(\gamma, \pi_0, \pi_1)$ that is independent of μ .*

Theorem 4 and Proposition 3 specialize immediately for the cases where heterogeneity is either only in beliefs or only in ambiguity aversion.

COROLLARY 1: *Assume exactly three distinct likelihood ratios. Then,*

- (a) *Polarization with Homogeneous Beliefs: Two individuals with beliefs η and constant relative ambiguity aversions $\hat{\gamma}$ and $\check{\gamma}$ exhibit polarization after observing the intermediate signal if and only if*

$$\tau(\hat{\gamma}, \pi_0, \pi_1) \leq \eta \leq \tau(\check{\gamma}, \pi_0, \pi_1)$$

with at least one inequality strict; and

- (b) *Polarization with Heterogeneous Beliefs: Two individuals with constant relative ambiguity aversion γ and beliefs $\hat{\eta}$ and $\check{\eta}$ exhibit polarization after observing the intermediate signal if and only if*

$$\hat{\eta} \geq \tau(\gamma, \pi_0, \pi_1) \geq \check{\eta}$$

with at least one inequality strict.

When the intermediate signal is a neutral signal, thresholds always lie strictly between 0 and 1 and take a particularly simple form:

COROLLARY 2: *Assume constant relative ambiguity aversion $\gamma > 0$ and exactly three distinct likelihood ratios. If x^M is a neutral signal, the threshold $\tau(\gamma, \pi_0, \pi_1)$ is*

$$\frac{1}{1 + \left(\frac{\pi_1(x^H)}{\pi_0(x^H)} \frac{\pi_1(x^L)}{\pi_0(x^L)} \right)^{\gamma + \frac{1}{2}}} \in (0, 1).$$

Heterogeneous tastes or beliefs are the source of polarization under ambiguity in Theorem 4. But this cannot explain the polarization observed by Darley and Gross (1983), where the groups exhibiting polarization were homogeneous. In their study, heterogeneity was induced across groups at the interim stage by showing them different initial evidence. Our previous results, together with their extension to prediction after observing multiple signals as described in the working paper version (Baliga, Hanany, and Klibanoff 2013), imply that exactly this device can generate polarization in an ex ante homogeneous prediction problem. For example, suppose there are two individuals with a common coefficient of relative ambiguity aversion $\gamma > 0$, a common prior $\mu = \frac{1}{2}$ and signals with three distinct likelihood ratios and symmetric likelihoods (i.e., $\pi_0(x^L) = \pi_1(x^H)$, $\pi_0(x^M) = \pi_1(x^M)$, and $\pi_0(x^H) = \pi_1(x^L)$). The individuals are allowed to condition their prediction on two conditionally independent observations. Suppose one individual observes the sequence $\{x^L, x^M\}$, while the other observes the sequence $\{x^H, x^M\}$. The multisignal version of Theorem 2 implies that after one observation the first individual will have updated beliefs $\check{\eta} < \frac{1}{2}$, and the second individual will have updated beliefs $\hat{\eta} > \frac{1}{2}$. From Corollary 2 and symmetry of the likelihoods, the threshold for updating upon observing x^M is $\tau(\gamma, \pi_0, \pi_1) = \frac{1}{2}$. Since the beliefs $\check{\eta}$ and $\hat{\eta}$ are on opposite sides of this threshold, Theorem 4 implies that polarization will occur after the second observation, x^M .

More generally, as long as the threshold is interior and enough observations are available, polarization is possible after an intermediate signal. This follows since if one individual observes a long sequence of high signals and another observes a long sequence of low signals, their posteriors will end up on different sides of this threshold. If they then observe a common intermediate signal, they will update in opposite directions and polarize. A formal statement of this result on polarization in a homogeneous environment may be found in Baliga, Hanany, and Klibanoff (2013, Theorem 3.4).

III. Related Literature and Concluding Remarks

A. Related Literature

Dixit and Weibull (2007) show that polarization cannot occur under Bayesian updating in the standard linear-normal model where individuals' (different) priors and (common) noise are normally distributed. Signals in this model satisfy the

monotone likelihood ratio property (MLRP). They also argue via example that polarization can occur if signals do not satisfy MLRP. On closer inspection, however, their examples violating MLRP do not display polarization. In fact, our Theorem 1 shows that polarization cannot occur under Bayesian updating, whether MLRP or normality holds or not. Instead, in their examples, while the means or the medians of two individuals' beliefs move further apart after observing a common signal, their beliefs are not further apart according to stochastic dominance.

Acemoglu, Chernozhukov, and Yildiz (2009) study asymptotic disagreement in a model where individuals have different priors on parameters and also different distributions on signals conditional on the parameter. They show that posteriors on parameters can diverge. Kondor (2012) shows that polarization can be generated when individuals see different private signals that are correlated with a common public signal. Andreoni and Mylovannov (2012) provide a similar theory and test their model experimentally. Rabin and Schrag (1999) study a model of confirmatory bias where agents ignore signals that do not conform with their first impressions, and thus updating is simply assumed to be biased in the direction of current beliefs, directly generating polarization. Notice that in all of these papers, individuals may disagree about the likelihood of a publicly observed signal conditional on the parameter. This is likely to be a common, if unsurprising, source of polarization. In contrast, in our model, conditional on the parameter, all individuals agree on the distribution over signals and their independence, and yet an interesting theory of polarization still emerges.

The only other paper we know of relating ambiguity to polarization is Zimmer and Ludwig (2009). They study particular forms of dynamically inconsistent updating in a model where agents are Choquet expected utility maximizers (Schmeidler 1989), and polarization is defined as divergence of "expected" signal probabilities as the number of observations goes to infinity. This contrasts with our model, where updating is optimal in the sense of dynamic consistency, beliefs have the standard additive form, and polarization is defined after any signal realization rather than as a limit phenomenon.

B. *Concluding Remarks*

The arrival of information changes the hedging motive of ambiguity-averse individuals. Optimal (i.e., dynamically consistent) updating must counteract this hedging effect in addition to the more familiar likelihood effect. We show that this delivers a theory of polarization—describing when it can occur and when it cannot.

The model and theory can be extended in several ways. First, we have assumed the individual has perfect foresight of the number of observations that will be available before he needs to take an action and that there is only one action required in the problem. Suppose instead that foresight is limited and the individual believes that he must take an action after fewer observations than will, in reality, be available. This is a natural description of the approach plausibly taken by subjects in the experiments of Darley and Gross (1983). Suppose (i) the individual uses dynamically consistent updating in the part of the problem he foresees; and (ii) when faced with the unforeseen continuation problem, he applies dynamically consistent updating to the continuation starting from the posterior beliefs inherited from the foreseen problem. Then, the possibility of polarization and the logic behind it described in our analysis continue to hold.

Second, our results are developed using the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005). One benefit of using this model is that, like the standard Bayesian expected utility model, it allows us to describe beliefs, and thus polarization, through a probability measure. It is not obvious how to best define polarization for other models of ambiguity-averse preferences. Putting that aside, the fundamental connection between ambiguity aversion and dynamically consistent updating that must counteract a hedging effect is present in any complete preference model of ambiguity aversion. Exactly when such a connection generates polarization is likely to vary with choice of model, and spelling out these conditions for different models is left for future research. Such an investigation may require additional tools, as the smooth ambiguity model allows us to characterize the unique solution of our prediction problem using first-order conditions.

Finally, we have assumed fully dynamically consistent updating. As was mentioned in the opening section, the effects we identify continue to generate polarization even under substantially weaker assumptions. For example, after observing a signal, suppose the individual maximizes a weighted sum of utility under dynamically consistent updating and utility under Bayesian updating.⁷ As long as there is strictly positive weight on the former, all of our qualitative results on polarization under ambiguity aversion are preserved.

REFERENCES

- Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz.** 2009. "Fragility of Asymptotic Agreement under Bayesian Learning." Unpublished.
- Andreoni, James, and Tymofiy Mylovanov.** 2012. "Diverging Opinions." *American Economic Journal: Microeconomics* 4 (1): 209–32.
- Baliga, Sandeep, Eran Hanany, and Peter Klibanoff.** 2013. "Polarization and Ambiguity." Center for Mathematical Studies in Economics and Management Science Discussion Paper 1558.
- Darley, John M., and Paget H. Gross.** 1983. "A Hypothesis-Confirming Bias in Labeling Effects." *Journal of Personality and Social Psychology* 44 (1): 20–33.
- Dixit, Avinash K., and Jörgen W. Weibull.** 2007. "Political Polarization." *Proceedings of the National Academy of Sciences* 104 (18): 7351–56.
- Gul, Faruk, and Wolfgang Pesendorfer.** 2001. "Temptation and Self-Control." *Econometrica* 69 (6): 1403–35.
- Hanany, Eran, and Peter Klibanoff.** 2007. "Updating Preferences with Multiple Priors." *Theoretical Economics* 2 (3): 261–98.
- Hanany, Eran, and Peter Klibanoff.** 2009. "Updating Ambiguity Averse Preferences." *B.E. Journal of Theoretical Economics: Advances* 9 (1): Article 37.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji.** 2005. "A Smooth Model of Decision Making under Ambiguity." *Econometrica* 73 (6): 1849–92.
- Kondor, Péter.** 2012. "The More We Know about the Fundamentals, the Less We Agree on the Price." *Review of Economic Studies* 79 (3): 1175–1207.
- Rabin, Matthew, and Joel L. Schrag.** 1999. "First Impressions Matter: A Model of Confirmatory Bias." *Quarterly Journal of Economics* 114 (1): 37–82.
- Schmeidler, David.** 1989. "Subjective Probability and Expected Utility without Additivity." *Econometrica* 57 (3): 571–87.
- Zimper, Alexander, and Alexander Ludwig.** 2009. "On Attitude Polarization under Bayesian Learning with Non-additive Beliefs." *Journal of Risk and Uncertainty* 39 (2): 181–212.

⁷One foundation for such a model is the temptation and costly self-control representation of Gul and Pesendorfer (2001) with the normative preference generated from dynamically consistent updating and the temptation preference generated from Bayesian updating.