Supplement to

Asymmetric Gain-Loss Reference Dependence and Attitudes towards Uncertainty

Juan Sebastián Lleras*, Evan Piermont†, Richard Svoboda‡

May 23, 2016

Abstract

This supplement to Lleras et al. [2016] examines equilibrium bidding behavior in first price sealed bid auctions when bidders have Asymmetric Gain-Loss Preferences. We find that equilibrium bidding strategies are higher than the corresponding strategies when bidders are risk-neutral. As such, the predictions are more closely aligned with experimental findings. In addition we examine qualitative and quantitative notions of comparative reference dependence.

1 Application to First Price Auctions

Experimental data shows systematic overbidding compared to the theoretical predictions of Nash Equilibrium with risk neutral bidders [Coppinger et al., 1980, Cox et al., 1988]. In this note, we aim to explain the empirically observed deviation from the Nash solution by appealing to bidders who have asymmetric gain-loss preferences. Lange and Ratan [2010] study the link between overbidding and reference dependence in greater depth than we do here. They examine an auction environment in which bidders have Kőzsegi and Rabin [2006] (KR) type reference dependence. However, to tractably tailor KR to the application they must make a number of simplifying assumptions.[1] Most importantly, they do not characterize the reference point as fixed point. Namely, they assume (as we do here), that for fixed bidding strategies of the other bidders, the reference point given

*jslleras@gmail.com

†ehp5@pitt.edu, University of Pittsburgh, Department of Economics

‡rts33@pitt.edu University of Pittsburgh, Department of Economics

[1] For specifics see the discussion in section 1.3 of Lleras et al. [2016]

bid b, is simply the expected utility of the auction given, b. So, in equilibrium, each bid is associated with a reference point and the collection of reference points affect the optimal bid, but which bid is optimal *does not* affect the individual reference points.

In spirit, Lange and Ratan [2010] are using AGL preferences (albeit in a more complicated multi-dimensional environment). Thus, it is not surprising that we find similar results. As such, this application is not so much intended to resolve the issue of overbidding in sealed bid auctions, as it is to reinforce the tractability and simplicity of the AGL representation when imported into larger models.

1.1 The Auction Environment

A single seller is selling an indivisible good. There are N potential buyers, represented by $i = \{1, \ldots, N\}$. Each buyer has a private value for the good, θ_i. Each θ_i is independently drawn from $[\underline{\theta}, \overline{\theta}] \subset \mathbb{R}^+$, with cumulative distribution function F. Denote the first derivative of the cumulative distribution function as f. The distribution is commonly known to all buyers. Furthermore, let $G(\theta)$ denote the probability of θ being the highest type out of N bidders, that is $F(\theta)^{N-1}$.

Bidders have AGL preferences over bid profiles, $b : [\underline{\theta}, \overline{\theta}]^N \to \mathbb{R}^N$, and a linear utility index $U_i(\cdot)$ defined over outcomes:

$$U_i(b) = \begin{cases} \theta_i - p_i(b) & \text{if the bidder wins the item and pays } p_i(b), \\ -p_i(b) & \text{if the bidder does not win the item and pays } p_i(b). \end{cases} \tag{1.1}$$

Note that without specifying the assignment rules of the auction, we do not have a full characterization of preference, as the mapping from bid profiles to equation (1.1) is undefined. Therefore we are abusing notation in assuming that U_i is defined over bid-functions, when it is actually defined over the composition of bid functions and, the as of yet unspecified, assignment rule. Given an assignment rule, each bid function constitutes a state contingent contract. In each state (a resolution of the type-space) the bidding profile specifies the bids, and the assignment rule specifies the outcomes to each bidder accordingly.

Ex-ante, an each bidder chooses as bid so as to maximize maximizes her expected AGL utility. The reference point, as discussed in the introduction of Lleras et al. [2016], is the expectation of this "first-order" utility with respect to the bidders' beliefs, μ_i. However, because we are interested pure strategy, monotone, symmetric Bayes-Nash equilibria, where the bidder with the highest valuation always wins the auction, we can assume without loss of generality that a bidder's belief that she is going to win is equivalent to her belief that she is the bidder with the highest value, as given by Bayes' rule and the commonly known distribution over types, F.

Therefore, the expected utility of bidder i can be written as:

$$\mathbb{E}_{F(\theta_{-i})}[U_i(b)] + \lambda \int_{\theta_{-i}:\theta_i < \max_j \theta_j} \big(\mathbb{E}_{F(\theta_{-i})}[U_i(b)] - U(b)\big) dF(\theta_{-i}). \tag{1.2}$$

This is the AGL-functional, defined in Lleras et al. [2016], and adapted to the auction setting. The first term is the expected utility and the second term is the expected gain-loss utility. Given the strategy of all other bidders, bidder i seeks to maximize (1.2) by choosing b_i (and hence choosing b subject to the restrictions of the other bidders).

Although a bidder's reference point is determined by the equilibrium strategies, the reference point itself is not an equilibrium condition (i.e., given fixed strategies of other bidders). Fixing the opponents strategies and the assignment rule, the choice of a bid is analogous to the choice from a set of acts. The acts pay according to the assignment rule and the probability of each state is determined by the strategies of other bidders, and the distribution over the type space. Each act carries with it a unique reference point (the expected first order utility) irrespective of which act is chosen.

As such we can employ the standard techniques in auction theory to solve for equilibrium bidding behavior. This, of course, would not be possible if the reference point was not uniquely identified, or depended on equilibrium conditions. Given the later, it would require that bidder i not only maximize her expected (reference dependent) utility subject to other bidders strategies, but also with respect to her own choice. In place of the relatively straightforward optimization problem, we would be instead tasked with finding a fixed point.

1.2 Equilibrium Bidding Functions

Assume that the seller implements a first price auction. Each buyer, i, makes a bid b_i. Then buyer $j = \underset{i}{\text{argmax}}\{b_i\}$ will receive the item, and pay her own bid, b_i; all other bidder pay nothing. Each bidder maximizes her expected utility, given by (1.2), according to her beliefs over other players strategies and types.

Theorem 1. *The unique symmetric monotone equilibrium is for a bidder of type θ to bid* *according to*

$$b^{AGL}(\theta) = \frac{\int_{\underline{\theta}}^{\theta} g(x)xdx}{G(\theta)\left[1+\lambda(1-G(\theta))\right]} + \frac{\int_{\underline{\theta}}^{\theta} g(x)x\left(\lambda\left(1-2G(x)\right)\right)dx}{G(\theta)\left[1+\lambda(1-G(\theta))\right]}.$$

Moreover, if $\lambda \in (-1,0)$ then there is a $\hat{\theta}$ such that for all $\theta \geq \hat{\theta}$, θ will overbid.

Denote the risk neutral SEU optimal bid of type θ as in a first price auction as $b^{RN}(\theta)$. Recall that symmetric bidding function is given by $b^{RN}(\theta) = \frac{\int_{\underline{\theta}}^{\theta} g(x)xdx}{G(\theta)}$. To see that there exists some threshold, above which all types will bid higher than the RNNE, note that when $\lambda \in (-1,0)$, the first term is larger than b^{RN} for all θ. Thus, all bidders for which the second term is positive will unambiguously bid larger than the RNNE. Basic continuity arguments suffice to show that this set exists and is nontrivial.

The effect here is due to the asymmetry (between gains and losses) of the preferences. It is the relatively large disutility from not meeting the reference point that motivates overbidding. Increasing the bid has two opposing effects on the reference point (and therefore on expected utility). First, as the bid increases, the expected payment is increasing. Thus, increasing the bid reduces profit. On the other hand, as the bid increases so does the bidder's probability of winning. Higher type bidders have a higher reference point, and so will have a tendency to overbid by more. The extra cost associated with the higher bid is, so to speak, payment to cover the now increased cost of losing.

We now turn to the mathematically simple environment where bidders types are uniformly distributed on the [0,1] interval. From these examples we can extract more intuition about the optimal bidding strategies. Additionally, most of the experimental evidence of overbidding arises from experiments where types are uniformly distributed.

With this distribution in mind, we can greatly simplify the bidding functions.

$$b^{RN}(\theta) = \frac{N-1}{N}\theta, \qquad \text{and,} \qquad b^{AGL}(\theta) = b^{RN}(\theta)\left[\frac{\left(1+\lambda\left[1-\frac{2N\theta^{N-1}}{(2N-1)}\right]\right)}{\left[1+\lambda(1-\theta^{N-1})\right]}\right].$$

From this we obtain the following result:

Corollary 2. *When the type space is uniformly distributed and $\lambda \in (-1,0)$, the optimal bid is weakly larger than $b^{RN}(\theta)$ for all types, and strictly larger for $\theta > 0$.*

Proof. This follows directly from setting $G(\theta) = \theta^{N-1}$, and verifying that the bracketed term in the simplified expression is always weakly greater than 1. ■

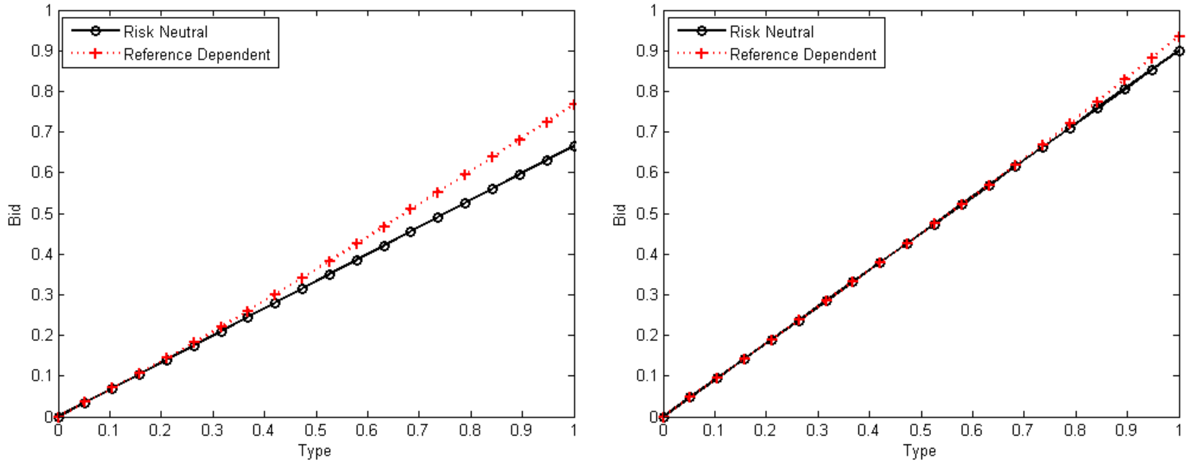

(a) $N = 3$.　　(b) $N = 10$.

Figure 1: Optimal Bids In a First Price Auction, $\lambda = -\frac{3}{4}$

By specifying the distribution, we can sign the difference in the optimal bid; AGL type preference can help to explain the persistence of overbidding in auctions. Figure 1 plots the optimal bids given $\lambda = -\frac{3}{4}$, in auctions with 3 and 10 bidders. It is evident from the graph that the magnitude by which a type overbids is monotonically increasing in type. Further, the overbidding diminishes as the number of bidders increases –the reference effect is smaller as the probability of winning (and so, the reference point) is decreasing in the number of bidders.

1.3 Proof of Theorem 1

We utilize the fact that we are searching for a symmetric and monotone equilibrium. Symmetry dictates the optimal bid function is invertible. Monotonicity dictates, the probability of bidder i winning the auction is $G(b^{-1}(b_i))$.

A loss-biased bidder will never bid above her type. This can be seen by noting that a bid, $b_i > \theta_i$ is strictly dominated by $b_i - \epsilon$ for some small enough ϵ. This is because lowering the bid reduces the likelihood of a losing state (winning the auction and paying more than her type –a strictly negative payoff) and weakly increase the payoff in all states.

Therefore, the reference point for any equilibrium bidding strategy is bounded weakly between 0 and the difference between the bidders type and her bid. Thus, we know states in which the bidder loses the auction are considered losses and states where the she wins the auction are considered gains. Using these facts we can write the bidder's maximization problem (of equation (1.2))

$$\max_{b_i} \; G(b^{-1}(b_i))\,(\theta_i - b_i) + \lambda\left(1 - G(b^{-1}(b_i))\right)\left(G(b^{-1}(b_i))\,(\theta_i - b_i)\right)$$

The first order condition is:

$$G(b^{-1}(b_i)) + G(b^{-1}(b_i))\lambda\left(1 - G(b^{-1}(b_i))\right) = \left(1 + \left[\lambda\left(1 - 2G(b^{-1}(b_i))\right)\right]\right)\frac{g(b^{-1}(b_i))}{b'(b^{-1}(b_i))}(\theta_i - b_i)$$

In equilibrium, we know that players will play monotonic and symmetric strategies. Therefore, each player's bid is a function of her type: $b_i = b(\theta_i)$. Further, $b^{-1}(b_i) = \theta_i$. Finally, we can drop the subscript i:

$$b'(\theta)G(\theta) + b'(\theta)G(\theta)\lambda\,(1 - G(\theta)) = (1 + [\lambda\,(1 - 2G(\theta))])\,g(\theta)\,(\theta - b(\theta))$$

Subtracting the last term from both sides and rearranging

$$\begin{gathered} b'(\theta)G(\theta) + b(\theta)g(\theta) + b'(\theta)G(\theta)\lambda\,(1 - G(\theta)) + b(\theta)\,[\lambda\,(1 - 2G(\theta))]\,g(\theta) \\ = \\ (1 + [\lambda\,(1 - 2G(\theta))])\,g(\theta)\theta \end{gathered}$$

We can rewrite the top line by noting that it is, in fact, a differential equation

$$\frac{d}{d\theta}\left[G(\theta)b(\theta) + G(\theta)b(\theta)(1 - G(\theta))\lambda\right] = (1 + [\lambda\,(1 - 2G(\theta))])\,g(\theta)\theta$$

Integrating both sides:

$$\int_{\underline{\theta}}^{\theta} \frac{d}{dx}\left[G(x)b(x) + G(x)b(x)(1 - G(x))\lambda\right]dx = \int_{\underline{\theta}}^{\theta} (1 + [\lambda\,(1 - 2G(x))])\,g(x)x\,dx$$

Using the boundary condition given that $G(\underline{\theta}) = 0$,

$$\left[G(\theta)b(\theta) + G(\theta)b(\theta)(1 - G(\theta))\right]\lambda = \int_{\underline{\theta}}^{\theta} (1 + [\lambda\,(1 - 2G(x))])\,g(x)x\,dx$$

Finally, we can solve for the bid function

$$b(\theta) = \frac{\int_{\underline{\theta}}^{\theta} (1 + [\lambda\,(1 - 2G(x))])\,g(x)x\,dx}{G(\theta)\,[1 + (1 - G(\theta))\lambda]}$$

We can rewrite this as the desired function.

$$b^I(\theta) = \frac{\int_{\underline{\theta}}^{\theta} g(x)x\,dx}{G(\theta)\,[1 + (1 - G(\theta))\lambda]} + \frac{\int_{\underline{\theta}}^{\theta} g(x)x\,(\lambda\,(1 - 2G(x)))\,dx}{G(\theta)\,[1 + (1 - G(\theta))\lambda]}$$

It remains to show that there is a $\hat{\theta}$ such that for all $\theta \geq \hat{\theta}$, θ will overbid. Since the first term of b^I is always weakly larger than b^{RN} is suffices to show that there is $\hat{\theta}$ such that for all $\theta \geq \hat{\theta}$

$$\frac{\int_{\underline{\theta}}^{\theta} g(x)x\,(\lambda\,(1 - 2G(x)))\,dx}{G(\theta)\,[1 + (1 - G(\theta))\lambda]} \geq 0 \tag{1.3}$$

We will show that for $\bar{\theta}$ the above holds strictly: the result then follows from the continuity of the bidding function. At $\bar{\theta}$, the denominator of (1.3) is 1 we need only show the numerator is strictly positive.

Through integration by parts, we obtain the following identity:

$$G^2(x)x|_{\underline{\theta}}^{\bar{\theta}} - \int_{\underline{\theta}}^{\bar{\theta}} G^2(x)dx - \int_{\underline{\theta}}^{\bar{\theta}} xg(x)G(x)dx = \int_{\underline{\theta}}^{\bar{\theta}} xg(x)G(x)dx$$

which allows us to rewrite the numerator of (1.3) as

$$\lambda\big[\int_{\underline{\theta}}^{\bar{\theta}} g(x)xdx - G^2(x)x|_{\underline{\theta}}^{\bar{\theta}} + \int_{\underline{\theta}}^{\bar{\theta}} G^2(x)dx\big]$$

Or, denoting μ_G as the expectation according to the distribution G,

$$\lambda\big[\mu_G - \bar{\theta} + \int_{\underline{\theta}}^{\bar{\theta}} G^2(x)dx\big]$$

And so, under the assumption that $\lambda < 0$, it suffices to show that

$$\int_{\underline{\theta}}^{\bar{\theta}} G^2(x)dx < \bar{\theta} - \mu_G \tag{1.4}$$

It is a well known consequence of Tonelli's Theorem that we can write the expectation of a non-negative random variable in terms of its CDF as

$$\mu_G = \int_0^{\infty} 1 - G(x)dx$$

which, given the support of $F(x)$, can be re-written as

$$\mu_G = \int_0^{\bar{\theta}} 1 - G(x)dx = \bar{\theta} - \int_{\underline{\theta}}^{\bar{\theta}} G(x)dx$$

So, since $G(\theta) \leq 1$ for all θ we know

$$\int_{\underline{\theta}}^{\bar{\theta}} G^2(x)dx < \int_{\underline{\theta}}^{\bar{\theta}} G(x)dx = \bar{\theta} - \mu_G$$

satisfying (1.4) and thus, proving the claim. ■

2 Comparative Statics

This section capitalizes on the connection to uncertainty attitudes in order to advance comparative statics results relating behavior to elements of the AGL representation. A natural measure for the degree and direction of reference effects is the gap between e_f and c_f, the hedge and the certainty equivalent. In the standard SEU model, e_f and c_f are the same, so the SEU model is the baseline case for reference effects.

Definition 1. Let $\succsim$ be a preference over $\mathcal{F}$. Say $\succsim$ is *gain-biased* if for all $f \in \mathcal{F}$, $c_f \succsim e_f$. Say $\succsim$ is *loss-biased* if for all $f \in \mathcal{F}$, $e_f \succsim c_f$.

Remark 1. An AGL decision maker is gain-biased (respectively, loss-biased) if and only if she is uncertainty seeking (resp., uncertainty averse) if and only if $\lambda \geq 0$. (resp., $\lambda \leq 0$).

Remark 1 follows immediately from the observation that $\mathbb{E}_{\mu}[f] = e_f$ and examination of the representing functionals.[2] This observation establishes a clear connection between the idea of "loss aversion" that has been prevalent since Prospect Theory, and uncertainty aversion.

If, however, we want to be able to compare two DM's *degree* of reference dependence we want to disentangle reference dependence from beliefs (which determine both c_f and e_f). To do this, we define $f \vee \bar{f}$, the join of a balanced pair $(f, \bar{f})$, as the act that gives the DM the best outcome between f and $\bar{f}$ for each $s \in S$.

Definition 2. Given any balanced pair $(f, \bar{f})$, define the act $f \vee \bar{f}$, the *join of* $(f, \bar{f})$ as

$$(f \vee \bar{f})(s) = \begin{cases} f(s) & \text{if} \quad f(s) \geq \bar{f}(s) \\ \bar{f}(s) & \text{if} \quad \bar{f}(s) > f(s) \end{cases} .$$

From the AGL representation, gain-loss utility depends on how much the act deviates state by state from e_f. $f \vee \bar{f}$ provides the absolute value of the state by state deviations of f from e_f. To determine the assessment of these deviations for each DM, we consider the hedge of the join, $e_{f \vee \bar{f}}$. Then, to capture reference-dependence behaviorally across DMs, we focus on acts that have the same hedge: if acts have different hedges, the reference effects can be confounded by the beliefs.

The intuition behind our comparative notion of "more loss biased" is that the DM prefers an act f with smaller deviations from e_f because the expected losses are smaller. Conversely, a DM with gain-bias prefers acts with larger deviations. Since we want to consider acts that have the same hedge, the comparative notions of "more gain-biased" and "more loss-biased" depends on a possibly different act for each DM: f for DM 1 and g for DM 2. Then the fact that $e_{f \vee \bar{f}}$ is larger than $e_{g \vee \bar{g}}$ provides a sufficient information to conclude that DM 1 is willing to pay more for f than DM 2 is willing to pay for g. We use the notation where e^i_f denotes the hedge of f for DM i.

Definition 3. Given two preference orders $\succsim_1$ and $\succsim_2$, say that $\succsim_1$ is *more gain-biased* than $\succsim_2$ (and $\succsim_2$ is more *loss-biased* than $\succsim_1$) if for any f, g with $e^1_f = e^2_g$ and $e^1_{f \vee \bar{f}} \geq e^2_{g \vee \bar{g}}$ for $i = 1, 2$, then for any $c \in \mathcal{F}_c$, $c \succsim_1 f$ implies $c \succsim_2 g$ and $c \succ_1 f$ implies $c \succ_2 g$.

The comparative notion of "more gain bias" or "more loss bias" is similar to the notion of uncertainty aversion in Ghirardato and Marinacci [2002]. For gain-loss preference the implication of comparative loss bias is the same as comparative uncertainty aversion, although in this case it holds only for acts that have the same hedge. This provides a behavioral way to compare reference effects across DMs.

[2] See Proposition A.5 in Lleras et al. [2016].

Theorem 3. *Let $\succsim_i$ admit an AGL representation given by (μ_i, λ_i) for $i = 1, 2$. Then $\succsim_1$ is more gain-biased than $\succsim_2$ if and only if $\lambda_1 \geq \lambda_2$.*

When $\mu_1 = \mu_2$, we can say more than Theorem 3. In this context, DM 1 being more gain biased is equivalent to $f \succsim_2 c$ implies $f \succsim_1 c$, for all acts f and constant acts c. Even further, if in addition both DM's are gain biased (or loss biased), then these equivalences can be extended to include $C_2 \subseteq C_1$ for the equivalent Maxmin/Maxmax representation from Theorem 3.1 in Lleras et al. [2016]. These addition equivalences stem from the fact that when DMs have the same belief, then for any $f \in \mathcal{F}$, $e^1_f = e^2_f$: loss bias is equivalent to the comparative notion of ambiguity aversion from Ghirardato and Marinacci [2002]. This further implies that whenever $\succsim_i$ is gain-biased or loss-biased for both DMs, the notion of loss bias is consistent with the representation of comparative ambiguity aversion derived from Gilboa and Schmeidler [1989] (that the more ambiguity averse DM should have a larger set of priors).

These comparative statics results establish an unexplored link between the absolute and comparative notions of gain or loss bias, and existing notions of uncertainty aversion which is worth further exploring. The initial motivation for studying uncertainty was due to the Ellsberg [1961] idea that DMs are not able to formulate unique probabilities over uncertain events. Many models with multiple priors have been developed to capture what is considered as "Ellsbergian behavior". Nonetheless, even if the DM is able to form a unique prior, having gain-loss considerations can appear to contaminate her prior in a way that gives rise to behavior embodied by some multiple priors model. Hence, for AGL preference a probabilistically sophisticated DM can appear to have multiple priors due to gain-loss asymmetry.

2.1 Proof of Theorem 3

Use the notation that superscripts denote the DM, e.g, e^i_f is the hedge of f for DM i.

(i) $\Rightarrow$ (ii). Let $\succsim_1$ be more gain-biased than $\succsim_2$. Consider any $f, g \in \mathcal{F}$ such that $e^1_f = e^2_g$ and $e^1_{f \vee \bar{f}} = e^2_{g \vee \bar{g}}$. So, by Proposition **??**, $\mathbb{E}_{\mu_1}[f \vee \bar{f}] = \mathbb{E}_{\mu_1}[g \vee \bar{g}]$. Now observe, by definition $f \vee \bar{f} = e_f + |f - e_f|$:

$$\mathbb{E}_{\mu_1}[e^1_f + |f - e^1_f|] = \mathbb{E}_{\mu_1}[e^2_g + |g - e^2_g|].$$

Since, $e^1_f = e^2_g$, this implies

$$\mathbb{E}_{\mu_1}[|f - e^1_f|] = \mathbb{E}_{\mu_2}[|g - e^2_g|]. \tag{2.1}$$

Suppose further, $g \succsim_2 c$ for any $c \in \mathcal{F}_c$, implies that $f \succsim_1 c$. Clearly, this is true if and only if $V_1(f) \geq V_2(g)$. We can write $V_1(f) \geq V_2(g)$ as defined in (**??**) as,

$$\mathbb{E}_{\mu_1}[f] + \frac{\lambda_1}{2}\mathbb{E}_{\mu_1}[|f - \mathbb{E}_{\mu_1}[f]|] \geq$$
$$\mathbb{E}_{\mu_2}[g] + \frac{\lambda_2}{2}\mathbb{E}_{\mu_2}[|g - \mathbb{E}_{\mu_2}[g]|].$$

Canceling according to (2.1) and $\mathbb{E}_{\mu_1}[f] = e^1_f = e^2_g = \mathbb{E}_{\mu_2}[g]$ we see $\lambda_1 \geq \lambda_2$.

(ii) $\Rightarrow$ (i). Let $\lambda_1 \geq \lambda_2$. Let $f, g \in \mathcal{F}$ be such that $e^1_f = e^2_g$, and $e^1_{f \vee \bar{f}} \geq e^2_{g \vee \bar{g}}$. Suppose for some $c \in \mathcal{F}_c$, $g \succsim_2 c$. Therefore, using the associate given by Proposition **??**,

$$V(g) = e^2_g + \frac{\lambda_2}{2}\mathbb{E}_{\mu_2}[|g - e^2_g]|] \geq c.$$

Since $e^1_{f \vee \bar{f}} \succsim_i e^2_{g \vee \bar{g}}$, $\mathbb{E}_{\mu_1}[|f - e^1_f|] = \mathbb{E}_{\mu_2}[|g - e^2_g|]$ by the same logic of (2.1). So,

$$\begin{aligned} V(f) &= e^1_f + \frac{\lambda_1}{2}\mathbb{E}_{\mu_1}[|f - e^1_f|] \\ &\geq e^2_g + \frac{\lambda_2}{2}\mathbb{E}_{\mu_2}[|g - e^2_g|] \\ &= V(g) \end{aligned}$$

Therefore $V(f) \geq c$, as desired. ■

References

V. M. Coppinger, V. L. Smith, and J. A. Titus. Incentives and behavior in english, dutch, and sealed-bid auctions. *Economic Inquiry*, 18(1):1–22, 1980.

J. C. Cox, V. L. Smith, and J. M. Walker. Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty*, 1(1):61–99, 1988.

D. Ellsberg. Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669, 1961.

P. Ghirardato and M. Marinacci. Ambiguity made precise: A comparative foundation. *Journal of Economic Theory*, 102(2):251–289, February 2002.

I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, April 1989.

B. Kőzsegi and M. Rabin. A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, November 2006.

A. Lange and A. Ratan. Multi-dimensional reference-dependent preferences in sealed-bid auctions–how (most) laboratory experiments differ from the field. *Games and Economic Behavior*, 68(2):634–645, 2010.

J. S. Lleras, E. Piermont, and R. Svoboda. Asymmetric gain-loss reference dependence and attitudes towards uncertainty. Mimeo, February 2016.