The Belief-Payoff Monotonicity Refinement

Andrew Kosenko*

November 11, 2022

Abstract

I define a novel equilibrium refinement for signaling games - *belief-payoff monotonicity*. I motivate this refinement, study its properties, argue that they are attractive, and relate it to existing refinements. Belief-payoff monotonicity stipulates that the receiver's beliefs upon observing an off-equilibrium path action are monotonic in the payoff increase for each type of sender from choosing such an action, under those beliefs. If multiple types benefit from a deviation, but their gains from that deviation are different, the receiver should assign higher probability to those types who benefit relatively more.

JEL Classification: D82, D83, C72.

Keywords: Equilibrium refinements, equilibrium selection, off-path beliefs, signaling games, cheap talk.

*Assistant Professor, School of Management, Marist College. 51 Fulton St, Poughkeepsie, NY, 12601. Email: kosenko.andrew@gmail.com. ORCID: 0000-0002-1227-7563. I am deeply grateful to Navin Kartik for his invaluable help and advice. I would also like to thank Yeon-Koo Che and Joseph Stiglitz for guidance and comments from which I have benefited immensely, as well as Ambuj Dewan, Nate Neligh, Teck Yong Tan, and the participants of the Columbia Microeconomic Theory Colloquium for discussions and input. Any errors my own.

1 Introduction

Signaling games are some of the most used theoretical representations of economic phenomena. Among the reasons they are appealing is their ability to capture significant parts of the economic environment (namely verbal communication where all messages are allowed, and the utility does not directly depend on the message) by incorporating private information in a tractable way. A signaling game is simply a game between two players (one is known as the leader, first mover, or *sender*, and the other the follower/second mover, or *receiver*) where one of the players - the sender - has many possible types which are known to her, but are not known by the receiver. The sender takes an action (sends a message), observed by all players, the receiver best-responds (given her beliefs about the sender's type, upon observing the message) by taking another action, and payoffs (functions of the type, the receiver's action, and perhaps the message) are realized.

Signaling games, for all their attractiveness, suffer from a defect - standard equilibrium concepts often do not always generate strong predictions in them; typically, there are many equilibria, of many kinds, and with many outcomes. The equilibria can be pooling (where all sender types take the same action), separating (where all sender types take different actions), or mixed/hybrid (where the actions taken by different sender types do not follow a simple pattern), or more frequently, of all three kinds. In other words, while signaling games are very useful representations, the predictive power of the typical solution concepts for them may be limited. One way of moving past this problem is to resort to so-called refinements of these equilibria to narrow down outcomes. A *refinement* is simply a condition on the equilibrium conditions; if an equilibrium does not satisfy such a condition it is said to fail the refinement. One then focuses only on the equilibria that survive the refinement as a way of strengthening the predictive content of the model.

The existing refinements aim to narrow down predictions by focusing on actions that are not taken on the equilibrium path of play. They rely on two principles; the first is often a version of the old adage "cui bono" - in other words, for which types is a particular action beneficial, relative to a particular equilibrium? The second principle seeks to adjust off-path beliefs of the receiver about the type of sender, following these off-path actions, so that they are consistent (in a sense appropriate to the setting) with the types who benefit from those actions. For example, if there is a single, unique type that benefits from a deviation, a widely used refinement, the "intuitive criterion", requires the receiver to believe with probability one that the deviation is coming from that type. There are a number of other refinements of this type, many (though not all) of them based on the concept of strategic stability proposed by Kohlberg and Mertens (1986). I review some of the relevant refinements below.[1]

I propose a new refinement, designed to work in a number of settings. I argue that it is not only a reasonable refinement, but that it is sometimes a necessary one. In addition, this new concept - termed *belief-payoff monotonicity*, or BPM for brevity - has a number of attractive properties. For example, it is strong in the sense that it can eliminate equilibria in some games where other refinements do not eliminate equilibria. Furthermore, it captures an appealing intuition - deviations must come from types that have the most to gain, if the receiver believes the "message" that is implicitly sent by such a deviation.[2]

The motivation for the refinement is this: suppose that there is an equilibrium and an associated (off-path) deviation so that multiple types benefit for some beliefs of the receiver, but that at least one type benefits relatively more than others. What should the receiver make of such a deviation, if observed? Certainly, any reasonable refinement would require the receiver to believe that the deviation is coming from the set of types that benefit, but are there any additional restrictions that may be desirable? Suppose for example, that while multiple types all benefit, one type benefits greatly, while others benefit only slightly; it is reasonable to stipulate that the receiver should believe that the deviation is coming from the type for whom the gain is greatest. It is precisely this intuition that BPM is trying to capture. This is also the reason for the nomenclature - the receiver's beliefs conditional upon an off-equilibrium path action should be monotonic in the payoff gain for each type of sender from choosing such an action. Thus, when multiple types benefit from a deviation, but their gains from that deviation are different, the receiver should assign higher probability to those types who benefit relatively more. The reader may also note that this is a *joint* type-message-belief condition.

There are several ideas at play here. The key ones are the idea of forward induction proposed by Kohlberg and Mertens (1986), and the notion of trembles introduced (al-

[1] I do not give definitions of these refinements, and instead point the reader to the original articles in the interests of keeping the present note short.

[2] Many other refinements attempt to capture a similar notion; I make these ideas precise and elucidate the ways in which out refinement is different in what follows.

beit in a slightly different setting - trembling-hand perfect equilibrium) by Selten (1965). Forward induction attempts to interpret deviations in some reasonable way - which is precisely what the BPM criterion is aiming to do by explicitly prescribing what the beliefs should be. Selten (1965) introduced the possibility that players may "tremble" and take non-best-response actions. Finally, Myerson (1978) proposed that if players do tremble, they should tremble lexicographically less often to actions that yield a lower payoff.[3] As discussed above, I adapt and unite these ideas and take the stand that deviations (interpreted as trembles) should be attributed to the types of sender than benefit the most from such a deviation, provided the receiver holds exactly the beliefs that make this true.

There are a few questions that are behind much of the reasoning on refinements and alternative equilibrium concepts - what *do* you make of a message that could have been sent, but wasn't, what *should* you make of it, and who would benefit as a result? The answer to these question is key in determining what sort of beliefs or equilibria are admissible; I explore a particular answer in this note. To keep the exposition brief, I rely on ideas that are standard in this literature; for this reason this note is perhaps most useful for readers already familiar with the literature on signaling games.

At this point the reader may justifiably wonder - why add a new refinement to the already large bestiary of such beasts? The reason is that this refinement turns out to work in a situation where others are unsatisfactory (see Kosenko (2022) for one such application). I view this refinement as not better or worse than others - but it may be helpful in *some* situations where others remain silent. In addition, this refinement is quite "strong" qualitatively in that if an equilibrium is ruled out by some other concept, it is likely[4] ruled out by BPM, so I view this refinement as one of last resort - if all others have failed, BPM may be a reasonable option.

2 Environment

There is a finite set of types for the sender: $\theta \in \Theta$, a finite set of states of the world $\omega \in \Omega$. Typically, the set of types of the sender and the set of states of the world are identified, but could, in principle, be different; in this short note I identify them for simplicity. Denote by

[3]Quantal response equilibrium of McKelvey and Palfrey (1995) captures a similar idea in experiments - players make mistakes with probability that is proportional to the loss of a particular action.

[4]In a sense discussed below.

$m \in M$ the message sent by the sender, and by $a \in A$ the action taken by the receiver. The utilities are $u^S(\theta, m, a)$ for the sender and $u^R(\theta, m, a)$ for the receiver. Denote by σ_θ^S and σ^R the respective strategies and let the final posterior beliefs of receiver be given by β.

Fix a perfect Bayesian equilibrium (PBE): $e = \{\sigma_\theta^S, \sigma^R, \beta\}$ with associated equilibrium utilities $u^{*S}(\theta)$; suppose for simplicity that A is a compact set and that $\beta \mapsto a(\beta)$ is one-to-one and onto; in particular this means that one can drop the a argument from the sender's utility. Say that e fails the criterion if there exists a type θ', a message[5] m', not sent in equilibrium e with positive probability, and a belief of the receiver $\beta(m')$ for which the following is true:

Definition 1 (Belief-Payoff Monotonicity Refinement - BPM)**.** *Let $e \triangleq \{\sigma_\theta^S, \sigma^R, \beta\}$ be an equilibrium and let $u^*(\theta)$ be the equilibrium utility of type θ. Define, for a fixed m, $\overline{u}(\theta_i) \triangleq \max_\beta u(m, \theta_i, \beta)$ and $\underline{u}(\theta_i) \triangleq \min_\beta u(m, \theta_i, \beta)$. An equilibrium is said fail the ϵ-BPM criterion if there is an experiment m, not chosen with positive probability in that equilibrium and a type of sender, θ_i, such that:*

i) Let $\beta \in \Delta(\Omega)$ be an arbitrary belief of the receiver and suppose that $\delta(m, \beta, \theta_i, e) \triangleq \frac{\hat{u}^S(m,\theta_i,\beta) - u^*(\theta_i)}{\overline{u}(\theta_i) - \underline{u}(\theta_i)} > 0$, for that belief.

ii) Denote by K be the set of types for which (i) is true; if K is empty BPM is inoperative so suppose that there is at least one type-message-belief triple for which i) holds. Let θ_i be the type for which the difference is greatest. If there is another type θ_j in K, for which $\delta(m, \beta, \theta_i, e) > \delta(m, \beta, \theta_j, e)$ then let $\beta(\theta_j|m) < \epsilon\beta(\theta_i|m)$, for some positive ϵ, with $\epsilon < \frac{1}{|K|}$. If there is yet another type θ_k such that $\delta(m, \beta, \theta_j, e) > \delta(m, \beta, \theta_k, e)$, then let $\beta(\theta_k|m) < \epsilon\beta(\theta_j|m)$, and so on.

iii) Beliefs are consistent: given the restrictions in (ii), the belief β is precisely the beliefs that makes (i) true.

The reason for the normalization in part i) of the definition is to make the definition stand up to affine transformations of the utility function (see also de Groot Ruiz et al. (2011)). The third part of the definition is a consistency requirement; it rules out situations such as the following. Suppose that the receiver believes that the deviation is coming from a particular type (say, type i), but it is type j that benefits more. Without the third

[5] I use the terminology of "messages" stemming from the cheap talk literature; this would just as well be some other "action".

requirement BPM would rule out such an equilibrium, but clearly beliefs in that case are not internally consistent or reasonable. Thus, one also has to check for internal consistency when applying BPM.

Say that an equilibrium fails the BPM criterion if it fails the ϵ-BPM criterion for every admissible ϵ with ϵ going to zero. However, I view ϵ-BPM as the more relevant refinement since it is more flexible[6]; the definition of BPM is stated as a limit since it is more intuitive and straightforward to apply.

This definition takes a clear, easily applicable stance on what beliefs should be off-path. There are, of course, other stipulations one can make; I now turn to these possibilites. One such stipulation is that the probability assigned to a deviation should be proportional to the gain for a type (so that, for example, if the gain for one type is twice the gain for another type, then the receiver should believe that the deviation is coming from the first type with probability two thirds, and from the second type with probability one third). This can be accommodated by choosing ϵ appropriately.

Another, perhaps more interesting issue is this: the definition given above fixes a belief, and then considers a particular deviation. However, given a belief, there may be multiple deviations for each type that can be beneficial - how should a sender "tremble" among them, and what should the receiver believe?

A reasonable and strong[7] definition may be the following. First, take an off-path belief for the receiver, and compute the relative utilities from deviating to all actions, for each type, given that belief. Then assume that each type will deviate to either sending the message that is most beneficial, or that each type will tremble among the possible messages that are beneficial, and that lower-gain messages will be sent with lower probability. And then apply ϵ-BPM for each message. This is arguably a more encompassing refinement. However, it is also more complex and makes even more assumptions about behavior; I thus focus on ϵ-BPM as a simpler and more easily applied definition.

Finally, I draw one useful connection between BPM and proper equilibrium (Myerson (1978)); both focus on similar trembles that are lexicographic in the (possible) gain. However, proper equilibrium requires one to assign smaller probabilities to *strategies* which are strictly dominated; whereas BPM requires the receiver to assign smaller probabilities

[6] In particular, in the typical case there may be multiple types that benefit from a deviation; the receiver may wish to assign some positive probability to the type that benefits less.

[7] And also related to reasoning behind proper equilibrium (Myerson (1978)).

to *types* that benefit relatively less.

3 Relationship to Other Refinements

3.1 Performance relative to stability-based refinements

In this section I explore the relationship of BPM to refinements that are based on the concept of strategic stability introduced by Kohlberg and Mertens (1986). These refinements operate by restricting off-path beliefs, as does BPM.

The first observation is that BPM is prior-independent (unlike, for instance, divine beliefs of Banks and Sobel (1987)). Furthermore, it can accommodate (i.e. make meaningful selections in) a version of cheap talk games. Generally speaking, criteria such as D1 do not have any bite in cheap talk games since they rely on messages that are unused in equilibrium, and in cheap talk games there is always an outcome-equivalent equilibrium in which all messages are used (for example, by randomizing over "unused" messages), and one is forced to resort to other equilibrium concepts (such as neologism-proofness that is discussed in the next section). BPM may, in fact, eliminate some cheap talk equilibria (as it does in the first chapter of this dissertation; see Kosenko (2022)).

Loosely speaking, BPM can be stronger or weaker than other concepts in the sense that it can do away with equilibria that are left untouched by other refinements, yet may also fail to eliminate other equilibria that are eliminated by other refinements in some cases.

I now turn to the question of examining the performance of BPM relative to other common refinements. Instead of formulating specific examples, I give simply a convenient representation of the relevant "moving parts" - the types of sender, the beliefs of the receiver, and the utility changes as functions of those beliefs. Well-chosen combinations of these moving parts will be sufficient to illustrate the main ideas. I illustrate the workings of BPM in relation to three commonly used (nested, and increasingly strict) refinements - the intuitive criterion (IC, Cho and Kreps (1987)), condition D1 (Banks and Sobel (1987)), and never a weak best response (NWBR, Cho and Sobel (1990)) criterion. There are many others in the same family (divinity, D2, iterated versions of these concepts, etc) but they are all nested in between these three, so by comparing BPM with them, I am also implicitly illustrating its potential relative to all the others.

To fix ideas, suppose for simplicity that there are only two types of sender - "red"

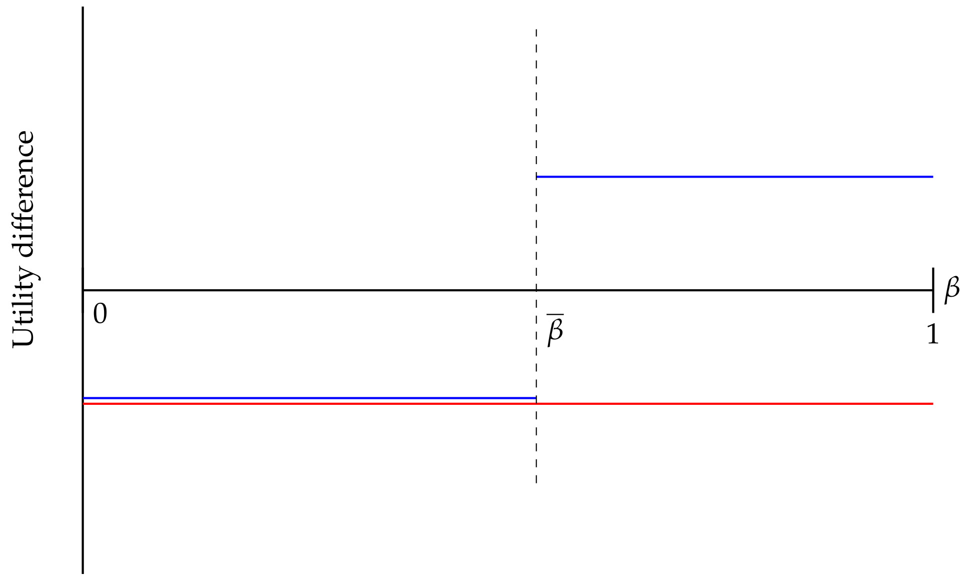

Figure 1: IC and BPM

and "blue", and fix some equilibrium as well as the corresponding equilibrium utilities. Suppose that the state of the world is the same as the type of the sender. Take a particular deviation, and consider the utilities of the two types as functions of the receiver's beliefs. Generically, the utility from a deviation will be different than the equilibrium utility; thus the relative utility *difference* from a deviation is plotted in the following figures.

Figure 1 illustrates how the intuitive criterion and BPM operate. In the typical case that is ruled out by IC, there are some beliefs of the receiver for which one type but not the other, benefits. More precisely, in figure 1, equilibrium is supported by beliefs $\beta \in [0, \overline{\beta})$, which make this deviation unattractive to either type. In that case, the equilibrium is said to fail IC - and it would also fail BPM, since the blue type has a profitable deviation that the red type does not: the set K from the BPM definition contains only the blue type for the relevant range of beliefs. In other words:

Observation 1. *Suppose that an equilibrium fails the intuitive criterion. Then it also fails the BPM criterion.*

I omit the proof for brevity, but the intuition is clear from figure 1 - if there is a unique type that benefits from a deviation for some beliefs, both concepts require one to believe that the deviation is coming from that type.

I now turn to the other frequently used refinement concept - condition D1 (Banks and Sobel (1987)) and show by example that BPM may or may not make the same equilibrium

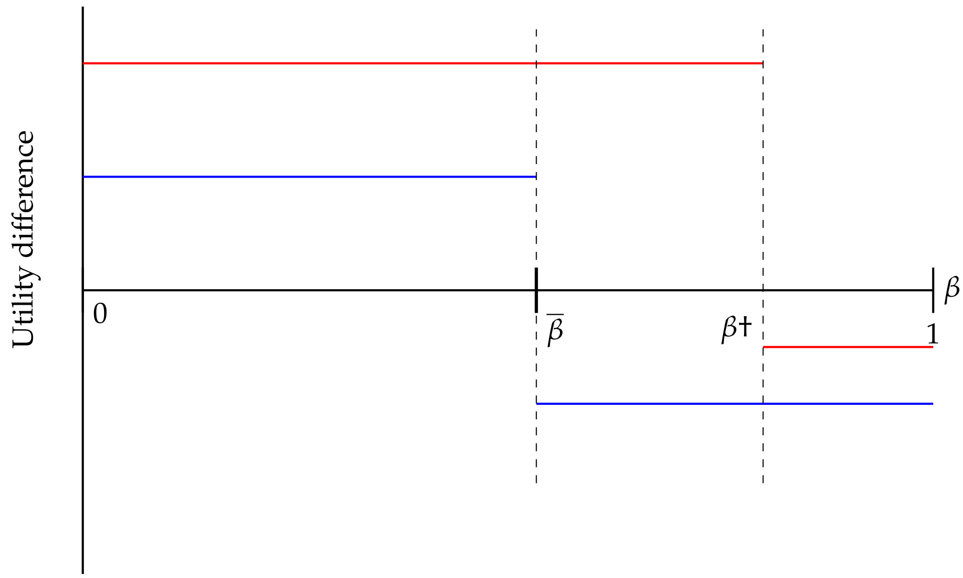

Figure 2: D1 and BPM make the same selection.

selection. First I examine a case where they do, this is illustrated in figure 2. D1 would eliminate this type-message pair (which clearly has to be supported by some belief $\beta \in (\beta\dagger, 1]$), since the set of beliefs for which the red type benefits ($[0, \beta\dagger)$) is a strict superset of the set of beliefs for which the blue type benefits ($[0, \overline{\beta})$). Similarly, BPM would eliminate this type-message pair since there are beliefs for which the red type benefits relatively more.

The never a weak best response (NWBR) criterion[8] is a strengthening of D1 that posits that whenever some type has a weak incentive to deviate (given some beliefs), then another type has a strict incentive to do so. A (perhaps typical) example is depicted in figure 3; NWBR would prune the blue type for this deviation since the red type has a strict incentive to deviate while the blue type is indifferent. BPM would do the same (for the same reason as in the IC example).

On the other hand, BPM may "disagree" with D1 - they may delete different type-message pairs. An example is shown in figure 4. Clearly, D1 would prune the blue type in this case, since the set of beliefs for which the red type benefits is strictly larger. However, for beliefs $\beta \in [0, \overline{\beta})$ it is the *blue* type that benefits more, and thus, BPM would delete the red type for those beliefs[9].

[8]This criterion is defined twice in the literature, once in the original Kohlberg and Mertens (1986) paper, and once in the Cho and Kreps (1987) work. The definitions are slightly different; I use the Cho-Kreps variant.

[9]Of course, for beliefs in $(\overline{\beta}, \beta\dagger)$ the two criteria would agree in deleting the blue type.

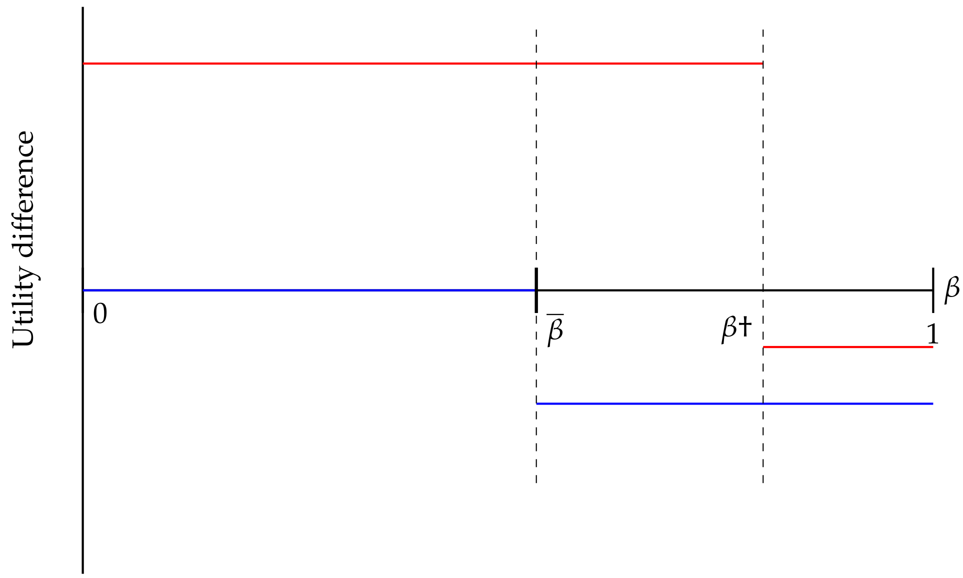

Figure 3: NWBR and BPM make the same selection.

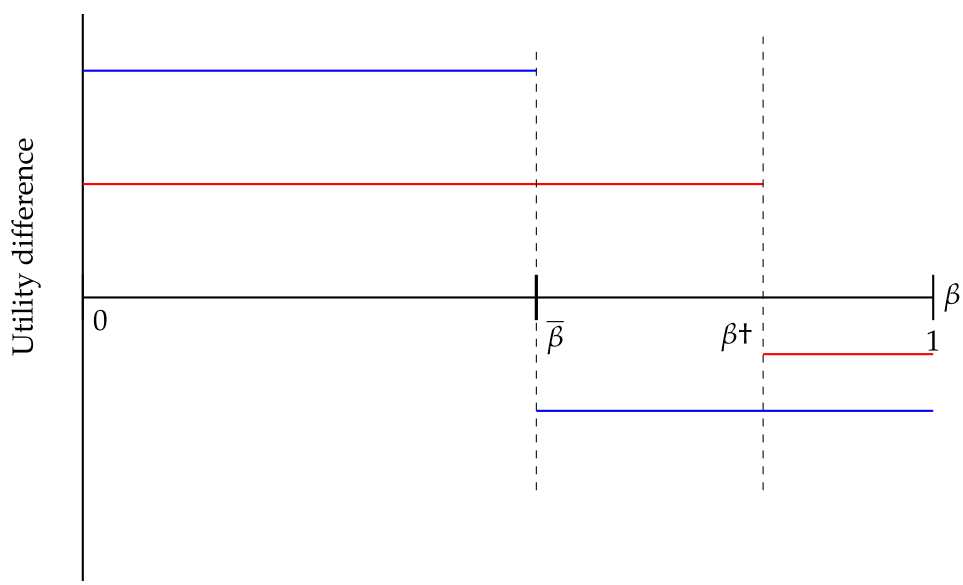

Figure 4: D1 and BPM make different selections.

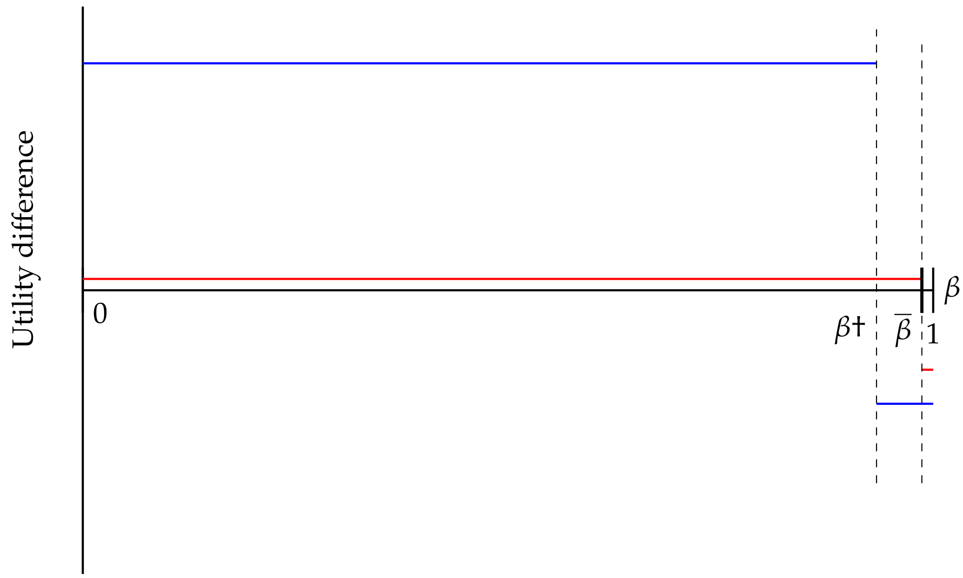

Figure 5: D1 vs. BPM: which is more convincing?

The two examples where D1 and BPM agree and disagree raise a reasonable question - which of the two refinements is more convincing? The figures also suggest that there is some interesting interplay between what D1 focuses on (the size of the set of beliefs for which a type benefits) and the magnitude of the gain from deviation, which is the focus of BPM. I illustrate this idea in figure 5 where as before, D1 and BPM would "disagree". However, depending on how one interprets trembles, either refinement may be more appealing. In this figure D1 would delete the blue type since the set of beliefs for which the red type benefits is larger. However, note that the red type benefits only a little (albeit for a strictly larger set of beliefs), while the blue type benefits quite a lot. In addition, the set of beliefs for which the red type benefits is not that much smaller than the corresponding set for the blue type. Given these two observations it is perfectly reasonable to delete the red type for this deviation, which is what BPM would prescribe. In short, this example shows that when BPM disagrees with other refinements, the question of which one is best depends on the particular case in point; either can be useful.

Finally, I give an example where D1 does not rule out any type-message pairs, while BPM does. In figure 6 condition D1 is inoperative since the relative sets are not nested. However, BPM would rule out both of these types.

The relation of BPM to stability-based refinement concepts is summarized in figure 7. The nested concepts are depicted in black circles (with inclusion representing subsump-

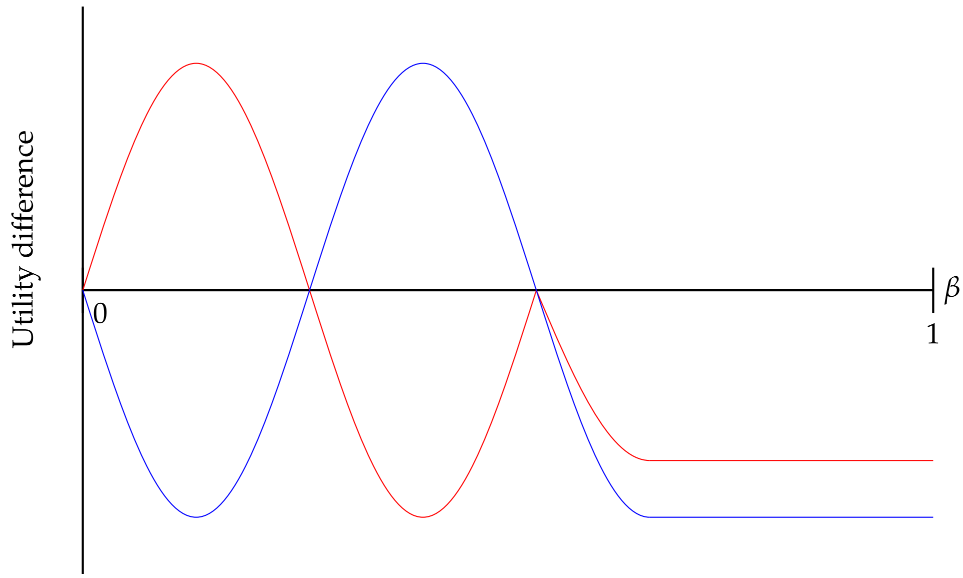

Figure 6: D1 does not apply, BPM does.

tion); the BPM refinement (represented by the red oval) may or may not agree with the refinements that are strictly stronger than IC (and it may, in fact, eliminate stable equilibria). However, whenever an equilibrium is eliminated by IC, it is also eliminated by BPM.

3.2 Performance relative to other refinements and equilibrium concepts in signaling games

Finally I turn to the question of the relationship between BPM and refinement concepts that are not based on the idea of strategic stability. One weakness of such refinements is that unlike stability based ones, these concepts often fail to exist.

For example, relative to the "money burning" idea introduced in Ben-Porath and Dekel (1992), BPM captures a similar idea. In "money burning" one can unilaterally "burn money" - destroy utility thus committing oneself to an action, which forces the other player to respond appropriately. The point is that with this possibility some equilibria are eliminated even without actually burning money on the equilibrium path - just the threat or possibility of this turns out to be enough. The high type of sender can "afford to burn" relatively more than the low type. In the absence of the option of burning actually payoffs (for example, in the standard examples from the Bayesian persuasion literature - an FDA drug trial and a court trial - it is not clear how one would go about burning util-

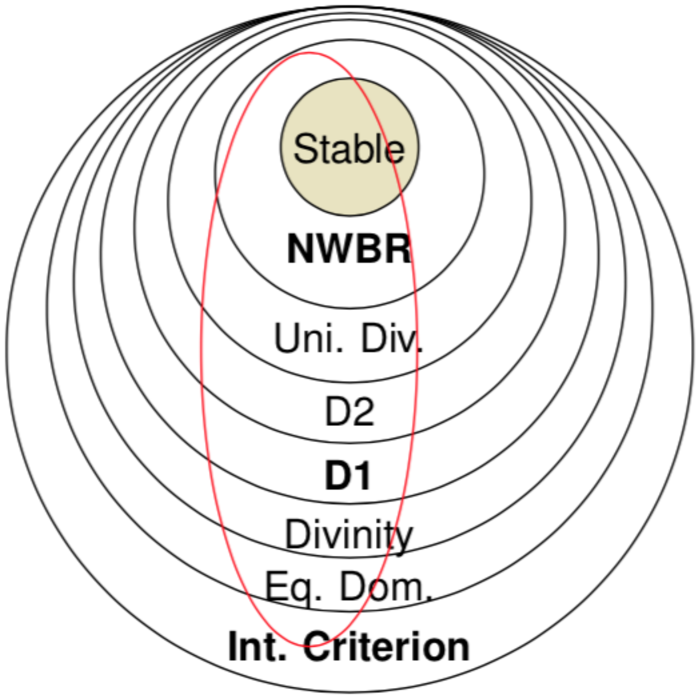

Figure 7: A bestiary of refinement concepts. BPM is depicted in red.

ity), ϵ-BPM offers a simple reduced-form definition that captures much of the same logic with similar results.

Similarly, relative to the concept of undefeated equilibria (Mailath et al. (1993)), BPM operates in much the same way. There is an example however (see Kosenko (2022)) where BPM rules out strictly more equilibria than undefeatedness. Like undefeated equilibrium, BPM may rule out *all* equilibria - i.e. it may fail to exist[10] However, Mailath et al. (1993) summarizes the undefeated equilibrium thus (p. 253):

> Consider a proposed sequential equilibrium and a message for player I that is not sent in equilibrium. Suppose there is an alternative sequential equilibrium in which some non-empty set of types of player I choose the given message and that that set is precisely the set of types who prefer the alternative equilibrium to the proposed equilibrium. The test requires player II's beliefs at that action in the original equilibrium to be consistent with this set. If beliefs are not consistent, say the second equilibrium *defeats* the proposed

[10] An example, unfortunately, is the standard purely dissipative Spencian signaling.

equilibrium.

Thus, Mailath et al. (1993) ask that there must be another *equilibrium* that defeats a putative equilibrium; BPM does not require that the alternative construction be an equilibrium to eliminate a putative equilibrium.

Similarly, the perfect sequential equilibria of Grossman and Perry (1986) tries to rationalize a deviation (once it occurs) by finding a set of types that benefit from such a deviation. They do so by defining a metastrategy that specifies how this is to be done; BPM would also eliminate equilibria that are not perfect sequential.

Note that both for perfect sequential and undefeated equilibria BPM would eliminate at least as many equilibria as either of these concepts. This is because if there exists an *equilibrium* that either defeats another, or a metastrategy that rationalizes a deviation, then surely there exist beliefs that satisfy the requirements for BPM to eliminate an equilibrium - simply use the type-message-beliefs triple in the defeating equilibrium.

Finally, BPM operates in a way that is analogous but not identical to the notion of neologism-proof equilibria of Farrell (1992). If an equilibrium is neologism-proof, it will survive BPM. However, BPM also takes a stand on how to "split" the probability weighting among the types in a self-signaling set; neologism proofness does not go that far. All three equilibrium concepts mentioned in this subsection may fail to exist, just like BPM.

4 Concluding Remarks

This note presents a definition a novel equilibrium refinement and explores its performance relative to other related concepts in the literature. The BPM criterion has some of the flavor of stability-based refinement, being a restriction on off-path beliefs, with the operative strength of other, newer equilibrium concepts. It appears to be stronger than most other refinements but suffers from lack of existence. It is presented as a "refinement of last resort" - this refinement may make a selection when other refinements fail to narrow down the set of possible outcomes.

References

[1] Banks, Jeffrey S. and Sobel, Joel, (1987). "Equilibrium Selection in Signaling Games." *Econometrica*, Econometric Society, vol. 55(3), pp. 647-61, May. https://doi.org/10.2307/1913604

[2] Ben-Porath, Elchanan and Eddie Dekel, (1992). "Signaling future actions and the potential for sacrifice." *Journal of Economic Theory*, Volume 57, Issue 1, Pages 36-51. https://doi.org/10.1016/S0022-0531(05)80039-0

[3] Blume, Andreas, (2012). "A class of strategy-correlated equilibria in sender-receiver games." *Games and Economic Behavior*, vol. 75. no. 2, pp. 510-517. https://doi.org/10.1016/j.geb.2012.03.008

[4] Blume, Andreas, Oliver J. Board and Kohei Kawamura, (2007). "Noisy Talk." *Theoretical Economics*, vol. 2, no. 4. pp.395-440.

[5] Chen, Y., Kartik, N. and Sobel, J., (2008). "Selecting Cheap-Talk Equilibria." *Econometrica*, vol. 76, pp. 117-136.

[6] Cho, In-Koo, and David M. Kreps, (1987). "Signaling Games and Stable Equilibria." *The Quarterly Journal of Economics*, vol. 102, no. 2, pp. 179-221. https://doi.org/10.2307/1885060

[7] Cho, In-Koo and Joel Sobel, (1990). "Strategic Stability and Uniqueness in Signaling Games". *Journal of Economic Theory*, 50, 381-413. https://doi.org/10.1016/0022-0531(90)90009-9

[8] de Groot Ruiz, Andrian, Theo Offerman and Sander Onderstal, (2013). "Equilibrium Selection in Cheap Talk Games: ACDC Rocks When Other Criteria Remain Silent." Working paper.

[9] Farrell, Joseph, (1993). "Meaning and Credibility in Cheap-Talk Games." *Games and Economic Behavior*, Volume 5, Issue 4, Pages 514-531. https://doi.org/10.3982/QE683

[10] Grossman, Sanford J. and Moty Perry, (1986). "Perfect Sequential Equilibrium." *Journal of Economic Theory*, Volume 39, Issue 1, Pages 97-119. https://doi.org/10.1016/0022-0531(86)90022-0

[11] Kohlberg, Elon, and Jean-Francois Mertens, (1986). "On the Strategic Stability of Equilibria." *Econometrica*, vol. 54, no. 5, pp. 1003-1037. https://doi.org/10.2307/1912320

[12] Kosenko, Andrew, (2022). "Constrained Persuasion with Private Information." Forthcoming in *The B.E. Journal of Theoretical Economics*. https://doi.org/10.1515/bejte-2021-0017

[13] Mailath, George J., Masahiro Okuno-Fujiwara, Andrew Postlewaite, (1993). "Belief-Based Refinements in Signalling Games." *Journal of Economic Theory*, Volume 60, Issue 2, Pages 241-276. https://doi.org/10.1006/jeth.1993.1043

[14] Selten, R. (1975). "Reexamination of the perfectness concept for equilibrium points in extensive games." *International Journal of Game Theory* 4, 25–55. https://doi.org/10.1007/BF01766400

[15] Mailath, G. J., Samuelson, L., and Swinkels, J. M.(1993). "Extensive Form Reasoning in Normal Form Games." *Econometrica*, vol. 61, no. 2, pp. 273-302. https://doi.org/10.2307/2951552

[16] Matthews, Steven A., Masahiro Okuno-Fujiwara, Andrew Postlewaite. (1991). "Refining cheap-talk equilibria." *Journal of Economic Theory*, Volume 55, Issue 2, Pages 247-273. https://doi.org/10.1016/0022-0531(91)90040-B

[17] McKelvey, R. D., and Palfrey, T. R, (1998). "Quantal Response Equilibrium for Extensive Form Games." *Experimental Economics*, 1, pp 9-41. https://doi.org/10.1023/A:1009905800005

[18] McKelvey, Richard D. and Thomas Palfrey. (1995). "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, 10(1), pp. 6-38. https://doi.org/10.1006/game.1995.1023

[19] Myerson, R, (1978). "Refinements of the Nash Equilibrium Concept." *International Journal of Game Theory*, vol. 7, pp. 73-80. https://doi.org/10.1007/BF01753236