A Nonatomic Game Involving Ambiguity

Jian Yang
Department of Management Science and Information Systems
Business School, Rutgers University
Newark, NJ 07102

March 2022

Essential Setups

- We study a *nonatomic game* (NG) with *ambiguity* considerations
- Besides a space I of *players* i and a space A of *actions* a, there is also a space Ω of *states* of the world ω
- Space R of *returns* r need not be real line $\Re$
- Under state ω, player i would receive return $\tilde{r}(\omega, i, a, \delta)$ when taking action a while other players form *external environment* δ
- For general *semi-anonymous* case, such external environments form space $\mathcal{D} \equiv \mathscr{P}(I \times A)$ of joint *player-action* distributions
- However, actual state ω is *not completely observable*

The Signal Space

- Instead, player i would receive a *signal* $\sigma \equiv \tilde{s}(\omega, i)$ under state ω
- We have let all players share a *common signal space* Σ which is conceptually *equivalent* to every player i having her own space $\Sigma'(i)$—just treat Σ as $\bigcup_{i \in I} \Sigma'(i)$
- Same signal σ could mean different things for different players; for instance, *color* red for player i and *sound* of wind for player i'
- When both state space Ω and signal space Σ are *finite*, there would be a finite space Σ^{Ω} of *signal vectors* $\vec{s} \equiv (s(\omega))_{\omega \in \Omega}$
- Each $\vec{s}$ would correspond to a subset $\mathcal{I}(\vec{s})$ of players i whose signals $\tilde{s}(\omega, i)$ are same as $s(\omega)$ prescribed by $\vec{s}$ for all states ω

State and Player Decompositions

- All players in same $\mathcal{I}(\vec{s})$ would for every signal σ, share a *common* subset $\mathcal{W}(\vec{s},\sigma)$ of states ω such that $s(\omega)=\sigma$
- Note $(\mathcal{W}(\vec{s},\sigma))_{\sigma\in\Sigma}$ would form a *decomposition of state space* Ω
- Also, $(\mathcal{I}(\vec{s}))_{\vec{s}\in\Sigma^\Omega}$ would form a *decomposition of player space* I
- We have an example with $|\Omega|=8$, $I=[0,1]$, and $|\Sigma|=3$
- A signal vector can be associated to an integer between 1 and $|\Sigma|^{|\Omega|}=3^8=6,561$; a number could be *skipped* if its corresponding $\mathcal{I}(\vec{s})$ happens to be *empty*

An Illustration with $|\Omega| = 8$, $I = [0, 1]$, and $|\Sigma| = 3$

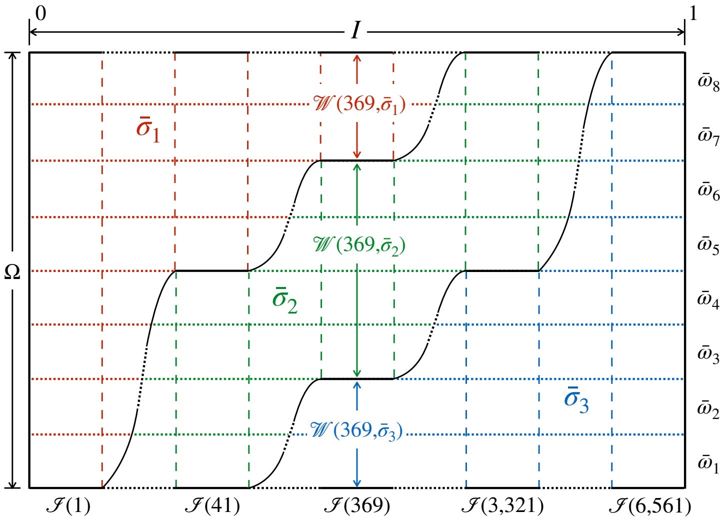

Return-distribution Vectors

- Call a player i who receives a signal σ the (i, σ)-player
- Under each state $\omega \in \mathcal{W}(\vec{s}, \sigma)$, this player could use her action lever to influence *distribution of return*, with deterministic return as a special case, that would come her way
- How each (i, σ)-player values all *return-distribution vectors* $\vec{\rho} \equiv (\rho(\omega))_{\omega \in \mathcal{W}(\vec{s}, \sigma)}$ or in parlance of Anscombe and Aumann (1963), *acts*, would be a key determinant of our game
- Under Harsanyi's (1967-8) *expected-utility Bayesian* framework, there would be a real-valued function $\tilde{u}(i, \sigma)$ and a prior $\tilde{k}(i, \sigma)$ such that on mind of (i, σ) is maximization of $\sum_{\omega \in \mathcal{W}(\vec{s}, \sigma)} \tilde{k}(\omega | i, \sigma) \cdot \int_R \tilde{u}(r | i, \sigma) \cdot [\rho(\omega)](dr)$

Risk and Ambiguity

- A *linear treatment* of return distributions can be legitimized by von Neumann and Morgenstein's (1944) axioms; use of a *single prior* can be couched on Savage's (1954) arguments
- Allais (1953) questioned whether people use *linear functionals* of return distributions to reach decisions
- Ellsberg (1961) argued that probabilities purportedly being assigned to different states of world are often *not known*
- For instance, there are probably not enough data to estimate chance for a new pandemic to occur in next two years
- Starting from Schmeidler (1989), researchers applied tools like Choquet integrations to *single-agent* decision making involving *general ambiguity attitudes*; see, e.g., Gilboa and Marinacci (2013)

Ambiguity Considerations

- Under axioms associated with *ambiguity aversion*, Gilboa and Schmeidler (1989) popularized *worst-prior* form in which our (i, σ)-player should maximize

$$\min_{k \in \tilde{K}(i,\sigma)} \left\{ \sum_{\omega \in \mathcal{W}(\vec{s},\sigma)} k(\omega) \cdot \int_R \tilde{u}(r|i,\sigma) \cdot [\rho(\omega)](dr) \right\}$$

  for some *ambiguity set* $\tilde{K}(i,\sigma)$ of priors k on $\mathcal{W}(\vec{s},\sigma)$

- We focus on *nonsingleton-prior* or more general ambiguity rather than *nonlinear-functional* risk attitudes

- Even when void of explicit risk considerations, most attention has been paid to ambiguity *aversion*

General Ambiguity Attitudes

- However, experiments involving human subjects showed ambiguity *seeking* could be equally prevalent; see, e.g., Curley and Yates (1989) and Charnes, Karni, and Levin (2013)
- We also believe optimistic assessment of uncertain gains is part of what drive people to participate in *auctions*, embark on exploratory *journeys*, and start new *companies*
- Given all these varieties of ambiguity scenarios, we believe it judicious to start from an all-inclusive *general case*
- Let us model each (i, σ)-player's *ambiguity attitude* by a *preference relationship* $\tilde{\psi}(i, \sigma)$ in space $(\mathscr{P}(R))^{\mathcal{W}(\vec{s},\sigma)}$ of *return-distribution vectors* $\vec{\rho} \equiv (\rho(\omega))_{\omega \in \mathcal{W}(\vec{s},\sigma)}$

More on Preferences

- A $\mathcal{W}(\vec{s},\sigma)$ might contain *summer* and *winter*, while our R might contain *ice cream* and *hot soup*
- A preference $\tilde{\psi}(i,\sigma)$ might spell out that
  $\vec{\rho}^{\,0} \equiv$"definitely ice cream in summer and definitely hot soup in winter" is *better than* $\vec{\rho}^{\,0.5} \equiv$"50% chance ice cream and 50% chance hot soup in either season"
  and that latter is *better than* $\vec{\rho}^{\,1} \equiv$"definitely hot soup in summer and definitely ice cream in winter"
- All $\tilde{\psi}(i,\sigma)$'s would constitute our *preference game*'s ambiguity profile $\tilde{\psi} \equiv (\tilde{\psi}(i,\sigma))_{\vec{s}\in\Sigma^{\Omega},i\in\mathcal{I}(\vec{s}),\sigma\in\Sigma}$

Game-theoretic Details

- When all players adopt a *strategy* $\mu \equiv (\mu(\cdot|i,\sigma))_{i \in I, \sigma \in \Sigma}$ where each $\mu(\cdot|i,\sigma)$ is an (i,σ)-dependent *distribution on actions* a, there would emerge for each state ω a *joint player-action distribution* $\Delta(\omega,\mu)$, which counts as *external environment* faced by all players under prevalent state ω and common strategy μ

- By taking any action a, a player i would reap $\tilde{r}(\omega, i, a, \Delta(\omega,\mu))$

- An (i,σ)-player with i in some $\mathcal{I}(\vec{s})$ would certainly want to take actions a so that resulting return-distribution vectors are *no more* $\tilde{\psi}(i,\sigma)$*-preferred to* by any other option

- There are *two* cases, one in which players can randomize among *actions* a that are no more preferred to, and another in which players can choose action *distributions* that are not preferred to

Two Equilibrium Notions

- The two possibilities would imply two equilibrium notions, namely, *action*- and *distribution*-based ones much as in a finite-player counterpart studied by Yang (2018)

- In option one, an *action*-based equilibrium μ would let almost every $\mu(\cdot|i,\sigma)$ devote entire weight to actions a that make $(\varepsilon(\tilde{r}(\omega,i,a,\Delta(\omega,\mu))))_{\omega\in\mathcal{W}(\vec{s},\sigma)}$ no more $\tilde{\psi}(i,\sigma)$-preferred to by any other *action* a', where $\varepsilon(r)$ stands for Dirac measure at r

- In option two, a *distribution*-based equilibrium μ would make each $(\sum_{a\in A}\mu(a|i,\sigma)\cdot\varepsilon(\tilde{r}(\omega,i,a,\Delta(\omega,\mu))))_{\omega\in\mathcal{W}(\vec{s},\sigma)}$ no more $\tilde{\psi}(i,\sigma)$-preferred to by any other *action distribution* α'

Equilibrium Existence and Continuity

- Action-based equilibria would always exist; further, set $\mathcal{E}^{\mathsf{a}}(\tilde{r}, \tilde{\psi})$ of such equilibria would be *upper hemi-continuous* in both *return function* $\tilde{r}$ and *preference profile* $\tilde{\psi}$
- We consider a preference ψ *convex* when $\vec{\rho}$ being no more ψ-preferred to than both $\vec{\rho}^{\,0}$ and $\vec{\rho}^{\,1}$ would together lead to $\vec{\rho}$ being no more ψ-preferred to than $(1-\beta)\vec{\rho}^{\,0} + \beta\vec{\rho}^{\,1}$ for any $\beta \in [0, 1]$
- When preferences $\tilde{\psi}(i, \sigma)$ are *convex*, set $\mathcal{E}^{\mathsf{d}}(\tilde{r}, \tilde{\psi})$ of distribution-based equilibria would be nonempty and upper hemi-continuous in $(\tilde{r}, \tilde{\psi})$
- We also get past this general *preference game* $\Gamma(\tilde{r}, \tilde{\psi})$ to some of its special cases that warrant separate attention

Special Cases

- When each preference $\tilde{\psi}(i,\sigma)$ is representable by a real-valued function $\tilde{\zeta}(\cdot|i,\sigma)$ on $|\mathcal{W}(\vec{s},\sigma)|$-long return-distribution vectors, we have a *satisfaction game* $\Gamma_{\mathsf{s}}(\tilde{r},\tilde{\zeta})$
- It would inherit properties from general case; especially, distribution-based equilibria would exist when each $\tilde{\zeta}(\cdot|i,\sigma)$ is *quasi-concave* to extent of $\tilde{\zeta}((1-\beta)\vec{\rho}^{\,0}+\beta\vec{\rho}^{\,1}|i,\sigma)$ being greater than $\tilde{\zeta}(\vec{\rho}^{\,0}|i,\sigma)\wedge\tilde{\zeta}(\vec{\rho}^{\,1}|i,\sigma)$ for any $\beta\in[0,1]$
- A further specialization would land us at an *alarmists' game* $\Gamma_{\mathsf{al}}(\tilde{r},\tilde{K},\tilde{u})$ when each satisfaction level $\tilde{\zeta}(\vec{\rho}|i,\sigma)$ at a given return-distribution vector $\vec{\rho}\equiv(\rho(\omega))_{\omega\in\mathcal{W}(\vec{s},\sigma)}$ is *worst* among an *ambiguity set* $\tilde{K}(i,\sigma)$ of priors $k\equiv(k(\omega))_{\omega\in\mathcal{W}(\vec{s},\sigma)}$ that contribute weights to *expected-utility* terms $\int_R\tilde{u}(r|i,\sigma)\cdot[\rho(\omega)](dr)$

Alarmists' and Enterprising Games

- Beyond *action-based* equilibria, an alarmists' game would always have *distribution-based* ones due to *concavity* of each $\tilde{\zeta}(\cdot|i,\sigma)$
- We also study opposite *enterprising game* $\Gamma_{\mathsf{en}}(\tilde{r},\tilde{K},\tilde{u})$ where *best* expected utility computed from among different prior choices in an ambiguity set is adopted
- Here, players exhibit *ambiguity-seeking* attitudes while betting optimistically on *most favorable* resolution of ambiguities
- Only action-based equilibria could be guaranteed for this case

The Expected-utility Bayesian Game

- On borderline between alarmists' and enterprising games lies *expected-utility Bayesian game* $\Gamma_{\text{bayes}}(\tilde{r}, \tilde{k}, \tilde{u})$ where earlier *prior sets* $\tilde{K}(i, \sigma)$ have degenerated into *single priors* $\tilde{k}(i, \sigma)$
- It would allow both *action*- and *distribution*-based equilibria
- Like Yang's (2018) finite counterparts, we examine *relationships* between the two equilibrium notions for *satisfaction games*
- A *distribution-based* equilibrium must be *action-based* for game such as an *enterprising* one with *convex* satisfaction functions
- The two equilibrium notions would *merge into one* at *expected-utility Bayesian game*

A Depiction of the Various Games

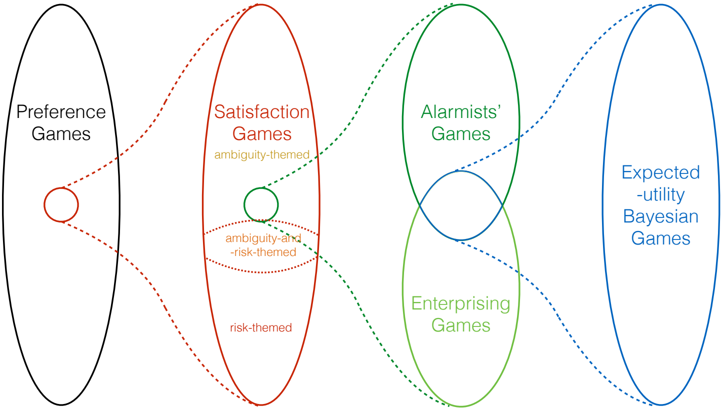

Connections with Finite Games

- Nonatomic games (NGs) are abstractions of real situations; knowledge on former ought to *help* us understand latter
- We can achieve something in vein of Yang's (2021) study of expected-utility Bayesian games
- A newly encountered subtlety is *divergent* behaviors of the two equilibrium notions
- Both action- and distribution-based NG equilibria could spur *mixed* correspondents that suffice as ϵ-*equilibria* in some probabilistic sense for *large enough finite games*
- Yet, only *action*-based NG equilibria could randomly generate ϵ-equilibrium *pure strategies* for large finite games

Representative Literature

- Normal-form NG: Schmeidler (1973), Mas-Colell (1984), Rath (1992), Balder (1995, 2002), Khan, Rath, and Sun (1997), Loeb and Sun (2006), Podczeck (2009), and Khan et al. (2013)
- Finite incomplete-information game: Harsanyi (1967-8), Radner and Rosenthal (1982), Milgrom and Weber (1985), Balder (1988), Kalai (2004), and He and Sun (2019)
- Bayesian NG: Khan and Rustichini (1991), Balder (1991), Balder and Rustichini (1994), Kim and Yannelis (1997), Carmona and Podczeck (2020), and Yang (2021)
- Finite game incorporating ambiguity: Dow and Werlang (1994), Klibanoff (1996), Lo (1998), Epstein (1997), Eichberger and Kelsey (2000), and Marinacci (2000)

Ambiguity on External Factors

- In most *ambiguity*-themed games, players were allowed to have qualms about *opponents' behaviors*
- Like finite counterpart Yang (2018), we focus on complementary situation where players have vagueness about *external factors*
- (i) mixed *strategies* chosen by players are often *enforceable*
  (ii) uncertainties about *state of world* can pose a much *bigger problem* than those about other players' behaviors
  (iii) *conventional tools* built on countably additive probabilities would suffice for our analyses
- Besides, uncertainty about opponents' *signals* would indirectly lead to uncertainty about their *preferences* as well as *behaviors*

Action- and Distribution-based Equilibria

- Two equilibrium concepts would stem from different ways of enforcing *mixed strategies*:

  action-based case—a player might be at almost total control of her own action in each play, but has to maintain agreed-upon *frequencies* to various actions in *long run*; this fits Dow and Werlang (1994) and Marinacci (2000)

  distribution-based case—a player might be given a *random number generator* whose output is private knowledge in-game but public knowledge post-game, and player has to abide by an agreed-upon *mapping* from random device's *output* to her *action*; this fits Klibanoff (1996) and Lo (1996)

- Kajii and Ui (2006) called first kind "equilibria in beliefs" and second kind "mixed equilibria"

Our Preference NG

- Our preference NG $\Gamma(\tilde{r}, \tilde{\psi})$ is built on
  a finite space Ω of *states*,
  space $I \equiv [0, 1]$ of *players*,
  a compact space R of *returns*,
  a finite space Σ of *signals*,
  a finite space A of *actions*,
  a state-player–to–signal *mapping* $\tilde{s}(\cdot, \cdot)$ that can be represented by player sets $\mathcal{I}(\vec{s})$ and state sets $\mathcal{W}(\vec{s}, \sigma)$,
  an *atomless player distribution* $\tilde{\lambda}$,
  a *return function* $\tilde{r}$ which contains elements $\tilde{r}(\omega, i, a, \delta)$, and
  a *preference profile* which contains elements $\tilde{\psi}(i, \sigma)$
- For each *irreflexive* and *transitive* preference $\tilde{\psi}(i, \sigma)$, $(\vec{\rho}, \vec{\rho}\,') \in \tilde{\psi}(i, \sigma)$ if and only if (i, σ)-player prefers $\vec{\rho} \equiv (\rho(\omega))_{\omega \in \mathcal{W}(\vec{s}, \sigma)}$ *no more than* $\vec{\rho}\,' \equiv (\rho'(\omega))_{\omega \in \mathcal{W}(\vec{s}, \sigma)}$

A Prevalent Environment

- In a particular state ω, a particular player i would receive a signal $\tilde{s}(\omega, i)$, prompting her to adopt an *action distribution* $\mu(i, \tilde{s}(\omega, i)) \equiv (\mu(a|i, \tilde{s}(\omega, i)))_{a \in A}$ under a given *strategy* μ
- In aggregate, *joint player-action distribution* felt by everyone would be $\Delta(\omega, \mu)$ such that for any player subset I' and action a,

$$[\Delta(\omega, \mu)] \left(I' \times \{a\} \right) \equiv \int_{I'} \mu(a|i, \tilde{s}(\omega, i)) \cdot \tilde{\lambda}(di)$$

- Mapping $\Delta(\omega, \cdot)$ from *strategies* to *externalities* under a given state ω would pose as an important feature for our game
- Note action distributions form simplex $\Theta_{|A|}$ in $[0, 1]^{|A|}$

Return Distributions and Their Vectors

- A player i's *action* distribution would drive her *return* distribution
- In any state ω, *return distribution* $\tilde{\rho}(\omega, i, \alpha, \mu|\tilde{r})$ achieved by her adopting action distribution α while all others adhering to strategy μ would satisfy, for any return subset R',

$$[\hat{\rho}(\omega, i, \alpha, \mu|\tilde{r})](R') = \sum_{a \in A} \alpha(a) \cdot [\varepsilon(\tilde{r}(\omega, i, a, \Delta(\omega, \mu)))](R')$$

- Given signal vector $\vec{s}$, a player $i \in \mathcal{I}(\vec{s})$, and a signal σ, *return-distribution vector* resulting from (i, σ)-player wielding an action distribution α while all others adopt a strategy μ would be

$$\overrightarrow{\rho}(i, \sigma, \alpha, \mu|\tilde{r}) \equiv (\hat{\rho}(\omega, i, \alpha, \mu|\tilde{r}))_{\omega \in \mathcal{W}(\vec{s}, \sigma)}$$

- For any action-distribution subset M', let us adopt notation

$$\overrightarrow{\rho}(i, \sigma, M', \mu|\tilde{r}) \equiv \{\overrightarrow{\rho}(i, \sigma, \alpha, \mu|\tilde{r}) : \alpha \in M'\}$$

Equilibrium-related Notions

- Given a background strategy μ, an (i,σ)-player's set $\mathbb{A}^{\mathsf{a}}(i,\sigma,\mu|\tilde{r},\tilde{\psi})$ of *best-responding pure actions* would be

$$\left\{a \in A:\ \overrightarrow{\rho}\,(i,\sigma,\varepsilon(A),\mu|\tilde{r}) \times \overrightarrow{\rho}\,(i,\sigma,\{\varepsilon(a)\},\mu|\tilde{r}) \subseteq \tilde{\psi}(i,\sigma)\right\},$$

  where $\varepsilon(A)$ should be understood as $\{\varepsilon(a):\ a \in A\}$

- Set $\mathbb{B}^{\mathsf{a}}(i,\sigma,\mu|\tilde{r},\tilde{\psi})$ of *best-responding action distributions* would be

$$\left\{\alpha \in \Theta_{|A|}:\ \sum_{a\in A} \alpha(a)\cdot[\varepsilon(a)]\left(\mathbb{A}^{\mathsf{a}}(i,\sigma,\mu|\tilde{r},\tilde{\psi})\right) = 1\right\}$$

- Set $\mathbb{B}^{\mathsf{d}}(i,\sigma,\mu|\tilde{r},\tilde{\psi})$ of *best-responding distributions* would be

$$\left\{\alpha \in \Theta_{|A|}:\ \overrightarrow{\rho}\,\left(i,\sigma,\Theta_{|A|},\mu|\tilde{r}\right) \times \overrightarrow{\rho}\,(i,\sigma,\{\alpha\},\mu|\tilde{r}) \subseteq \tilde{\psi}(i,\sigma)\right\}$$

Action- and Distribution-based Equilibria

- Define *action(distribution)-based best-responding correspondence*:

$$\mathbb{M}^{\mathsf{a}(\mathsf{d})}(\mu|\tilde{r},\tilde{\psi}) \equiv \{\mu' \in \mathcal{M} : \ \mu'(i,\sigma) \in \mathbb{B}^{\mathsf{a}(\mathsf{d})}(i,\sigma,\mu|\tilde{r},\tilde{\psi}), \text{ for any } \vec{s} \in \Sigma^{\Omega}, i \in \mathcal{I}(\vec{s}), \text{ and } \sigma\}$$

- We consider a strategy μ as belonging to set of *action(distribution)*-based *equilibria* $\mathcal{E}^{\mathsf{a}(\mathsf{d})}(\tilde{r},\tilde{\psi})$ if and only if it is a *fixed point* for $\mathbb{M}^{\mathsf{a}(\mathsf{d})}(\cdot|\tilde{r},\tilde{\psi})$; that is, $\mu \in \mathbb{M}^{\mathsf{a}(\mathsf{d})}(\mu|\tilde{r},\tilde{\psi})$

- Using approaches that ultimately tap into Fan-Glicksberg theorem, we can show that $\mathcal{E}^{\mathsf{a}}(\tilde{r},\tilde{\psi}) \neq \emptyset$ and that $\mathcal{E}^{\mathsf{a}}(\cdot,\cdot)$ is an *upper hemi-continuous* correspondence defined on space of *return functions* and *preference profiles*; same can be achieved for $\mathcal{E}^{\mathsf{d}}(\cdot,\cdot)$ when preferences are *convex* (mixture being innocuous)

Special Cases

- In more special *satisfaction game* $\Gamma_{\mathsf{s}}(\tilde{r}, \tilde{\zeta})$, an (i, σ)-player would either randomize over actions a whose $\varepsilon(a)$ maximize $\tilde{\zeta}\left(\overrightarrow{\rho}(i, \sigma, \cdot, \mu|\tilde{r})|i, \sigma\right)$ or adopt a distribution α that achieve same
- For even more special *alarmists' game* $\Gamma_{\mathsf{al}}(\tilde{r}, \tilde{K}, \tilde{u})$, to be maximized by each (i, σ)-player with $i \in \mathcal{I}(\vec{s})$ would be

$$\min_{k \in \tilde{K}(i,\sigma)} \left\{ \sum_{\omega \in \mathcal{W}(\vec{s},\sigma)} k(\omega) \cdot \int_R \tilde{u}(r|i,\sigma) \cdot \left[\hat{\rho}(\omega, i, \cdot, \mu|\tilde{r})\right](dr) \right\}$$

  Above *min* would become *max* in *enterprising game* $\Gamma_{\mathsf{en}}(\tilde{r}, \tilde{K}, \tilde{u})$

- For *expected-utility Bayesian game* $\Gamma_{\mathsf{bayes}}(\tilde{r}, \tilde{k}, \tilde{u})$, above would be

$$\sum_{\omega \in \mathcal{W}(\vec{s},\sigma)} \tilde{k}(\omega|i,\sigma) \cdot \int_R \tilde{u}(r|i,\sigma) \cdot \left[\hat{\rho}(\omega, i, \cdot, \mu|\tilde{r})\right](dr)$$

Relationships between Two Notions

- Due to *distribution*-based equilibria's "higher" requirements of competing with and winning over other *distributions* (with greater cardinalities) rather than *action-based* ones which deal with other *actions*, general message is that former are *rarer* than latter
- We need *convexity* of preferences and equivalently *quasi-concavity* of satisfaction functions for existence of *distribution*-based equilibria; there is no general guarantee for *enterprising game*
- For *satisfaction game* $\Gamma_{\mathsf{s}}(\tilde{r}, \tilde{\zeta})$, we can actually show that $\mathcal{E}^{\mathsf{d}}_{\mathsf{s}}(\tilde{r}, \tilde{\zeta}) \subseteq \mathcal{E}^{\mathsf{a}}_{\mathsf{s}}(\tilde{r}, \tilde{\zeta})$ when $\tilde{\zeta}(\cdot|i, \sigma)$'s are *convex*, leading for *enterprising game* to satisfy $\mathcal{E}^{\mathsf{d}}_{\mathsf{en}}(\tilde{r}, \tilde{K}, \tilde{u}) \subseteq \mathcal{E}^{\mathsf{a}}_{\mathsf{en}}(\tilde{r}, \tilde{K}, \tilde{u}) \neq \emptyset$
- We indeed have $\mathcal{E}^{\mathsf{d}}_{\mathsf{bayes}}(\tilde{r}, \tilde{k}, \tilde{u}) = \mathcal{E}^{\mathsf{a}}_{\mathsf{bayes}}(\tilde{r}, \tilde{k}, \tilde{u}) \neq \emptyset$

Finite n-player Games

- In an n-player game, player profile $i_{[n]} \equiv (i_m)_{m=1,\ldots,n}$ would be *randomly sampled* from distribution $\tilde{\lambda}$

- Rather than common $\Delta(\omega, \mu)$, *external environment* faced by player m would be empirical *player-action* distribution $\varepsilon(i_{[n],-m}, a_{[n],-m})$ which assigns a $(1/(n-1))$-weight to each (i_l, a_l)-realization for $l = 1, \ldots, m-1, m+1, \ldots, n$

- With $\overrightarrow{\rho}_n \left(i_m, \sigma_m, \alpha_m, \mu_{[n],-m}\right)$ representing corresponding *return-distribution vector*, an *action*-based ϵ-equilibrium would randomize over actions a_m that push

$$\overrightarrow{\rho}_n \left(i_m, \sigma_m, \varepsilon(A), \mu_{[n],-m}\right) \times \overrightarrow{\rho}_n \left(i_m, \sigma_m, \{\varepsilon(a_m)\}, \mu_{[n],-m}\right)$$

inside $(\tilde{\psi}(i_m, \sigma_m))^{\epsilon}$; *distribution*-based case would be analogous

Convergence Results

- With empirical player-action distribution converging to $\Delta(\omega, \mu)$ in some *probabilistic* sense when *player number* n tends to $+\infty$, we can show usefulness of NG *equilibria* in both *action*- and *distribution*-based senses in finite games

- For any NG equilibrium μ in *a(d)*-sense and $\epsilon > 0$,

$$\lim_{n \longrightarrow +\infty} \tilde{\lambda}^n(\{i_{[n]} \in I^n : \ \mu \text{ induces } \epsilon\text{-equilibrium in a(d)-sense} \\ \text{for } n\text{-player game with player profile } i_{[n]}\}) = 1$$

- After sampling over $\tilde{\lambda}$ to obtain *player profile* $i_{[n]}$, we can further sample over *strategy* $\mu(\cdot|i_m, \sigma)$ for each player m and each potential signal σ to obtain signal-based *pure-action* profile $\tilde{a}_{[n]} \equiv (\tilde{a}_m)_{m=1,\ldots,n} \equiv (\tilde{a}_m(\sigma))_{m=1,\ldots,n,\sigma\in\Sigma}$

A Mixed-to-pure Link

- We can construct a *strategy-dependent* probability $\Lambda(\mu)$ on product player–action-plan space $I \times A^{\Sigma}$ so that for any player subset I' and signal-based action plan $\tilde{a} \equiv (\tilde{a}(\sigma))_{\sigma \in \Sigma}$,

$$[\Lambda(\mu)]\,(I' \times \{\tilde{a}\}) \equiv \int_{I'} \left[\prod_{\sigma \in \Sigma} \mu(\tilde{a}(\sigma)|i,\sigma)\right] \cdot \tilde{\lambda}(di)$$

- For any NG equilibrium μ in *a*-sense and $\epsilon > 0$,

$$\lim_{n \longrightarrow +\infty} (\Lambda(\mu))^n \left(\{(i_{[n]}, \tilde{a}_{[n]}) \in (I \times A^{\Sigma})^n : \ \tilde{a}_{[n]} \text{ achieves } \epsilon\text{-equilibrium for } n\text{-player game with player profile } i_{[n]}\}\right) = 1$$

- A *distribution*-based counterpart seems *unlikely* because most any $\overrightarrow{\rho}(i, \sigma, \alpha, \mu)$ would be *inimitable* by any $\overrightarrow{\rho}_n(i, \sigma, \epsilon(a), \nu_{[n],-m})$ for some *pure* action a and strategy profile $\nu_{[n],-m}$

Concluding Remarks

- More general *state* space Ω, *signal* space Σ, and *action* space A can be considered by future research
- Removal or relaxation of certain *compactness* and *continuity* requirements should also be attempted
- More game varieties and more relationships between equilibrium notions await further exploration