

Mean Field Games of Timing and Models for Bank Runs

René Carmona¹ · François Delarue² ·
Daniel Lacker³

Published online: 30 June 2017
© Springer Science+Business Media, LLC 2017

Abstract The goal of the paper is to introduce a set of problems which we call mean field games of timing. We motivate the formulation by a dynamic model of bank run in a continuous-time setting. We briefly review the economic and game theoretic contributions at the root of our effort, and we develop a mathematical theory for continuous-time stochastic games where the strategic decisions of the players are merely choices of times at which they leave the game, and the interaction between the strategic players is of a mean field nature.

Keywords Mean field games · Bank runs · Optimal stopping · Supermodular games

1 Introduction

Our starting point is the set of early game theoretic models for the banking system due to Bryant [9] and Diamond and Dybvig [14] whose fundamental papers proposed banking models of a game played by depositors in which there always exist at least a good equilibrium and a bad equilibrium. Many generalizations followed, for example to include illiquidity effects and more random factors, and extended the scope of the

✉ Daniel Lacker
daniel_lacker@brown.edu

René Carmona
rcarmona@princeton.edu

François Delarue
francois.delarue@unice.fr

¹ Princeton University, Princeton, NJ, USA

² Université de Nice, Nice, France

³ Brown University, Providence, RI, USA

models beyond bank runs and deposit insurance to include financial intermediation, as in the work [33] of Rochet and Vives. There, the authors use the methodology of global games proposed by Morris and Shin [31] and the differences in opinions among investors to prove existence and uniqueness of a Nash equilibrium. They go on to analyze the economic and financial underpinnings of bank runs and propose a benchmark for the role of lenders of last resort. While still in a static framework, the work [18] of Green and Lin discusses stochastic equilibria (a.k.a. aggregate uncertainty) in a context which is very close to our notion of weak equilibrium, to be defined later in the paper.

Authors of the early game theoretic papers on bank runs quickly realized that their models exhibited what is now known as a *complementarity property*. Typically, if more depositors withdraw their funds early, then the probability of failure of the bank increases, and this further incentivizes early withdrawal. Mathematically, the eventual payoff to one depositor displays *increasing differences* with respect to the actions of the others depositors. This property is known as complementarity, and games with this property are called supermodular games. The equilibrium theory of these games hinges on their order structure more than their analytic properties (see, for example [20, 30]), using machinery first developed by Topkis [35, 36] and later refined by Milgrom and Roberts [30] and Vives [37].

A common feature of many bank run models is the symmetric or mean-field nature of the interaction between the depositors, and the goal of our paper is to take advantage of this property to develop a general mathematical theory. While most of the works cited above are static in nature, our interest in dynamic models of bank runs was sparked by a lecture of Olivier Gossner at a PIMS Workshop on Systemic Risk in July 2014 who attempted to extend to a continuous-time setting an earlier work of Rochet and Vives [33]. In this model, the common source of randomness comes from the value of the investments of the bank and the possible need for fire sales to face fund redemption, while the differences in the private signals of the investors contribute to the idiosyncratic sources of noise, ruling out undesirable equilibria. Another continuous-time bank run model worthy of mention can be found in the paper [19] by He and Xiong where the source of randomness comes from the staggered nature of the debt maturities.

With these bank run models in mind, we propose a general class of continuous-time models we call *mean field games of timing*, in which a continuum of agents strategically choose stopping times, i.e., times at which to exit the game. We present two different sets of results for two different regimes. Under the aforementioned complementarity property, we prove that “mean field equilibria” (MFE) exist and illustrate how to use them to construct approximate equilibria for the corresponding n -player games, and this is done for very general partial information structures. On the other hand, without complementarities, we derive an existence result for “weak MFE” under stronger continuity assumptions, and only in the full-information setting. We then connect weak MFE to n -player games by proving two modes of convergence. On the one hand, the equilibria themselves in the n -player game (if they exist) converge to weak MFE as $n \rightarrow \infty$. On the other hand, a weak MFE can be used to construct approximate equilibria for the n -player games.

Our models are closely related to the mean field games introduced independently by Lasry and Lions [29] and Caines et al. [21]. However in our models, agents act by choosing stopping times as opposed to control processes. We adhere to a purely proba-

bilistic approach, though in principle a PDE formulation is possible involving a variational inequality or free-boundary problem. Probabilistic methods in mean field game theory originated in [10], although our techniques are more closely related to the weak convergence and compactness arguments of [12, 26, 28]. While most (continuous-time) mean field game models involve agents choosing control processes as opposed to stopping times, a notable exception is the recent work of Nutz [32], which studies a tractable yet versatile model for which equilibria can be computed or at least characterized quite explicitly. Section 3.4 shows how this model fits into our framework.

Our existence result (Theorem 3.5) based on monotonicity properties resembles some recent papers on games with both complementarities and a continuum of agents. For instance, Adlakha and Johari [3] employ some similar techniques to study a discrete time mean field game with strategic complementarities. The work of Balbus et al. [5] on static games is also quite relevant, and it even includes a discussion of discrete-time “optimal stopping games,” although stochastic factors are absent from their model. See also [38] and its correction in [6] for related work on nonatomic supermodular games. The reader interested in games with complementarities may also consult the recent work of Acemoglu and Jensen [1, 2] on *aggregate games*, which closely resemble mean field games.

The technical crux of our proofs requires some new results, interesting in their own right, on progressive enlargements of filtrations [8, 24], particularly related to the “compatibility” or “immersion” property (also known as the H-hypothesis) which has recently seen renewed interest in light of its many applications in credit risk models. Our work necessitates a new characterization of when a filtration enlarged progressively by a random time satisfies this compatibility property: roughly speaking, if a filtration $\mathbb{F}$ is generated by a Wiener process W , and if it is enlarged progressively to $\mathbb{G}$ in the minimal way to render a given a random time τ a stopping time, then $\mathbb{F}$ is “compatible” with $\mathbb{G}$ if and only if there exists a sequence of $\mathbb{F}$ -stopping times τ_n such that (W, τ_n) converges to (W, τ) in distribution. This notion of compatibility arises naturally because of the central role played by weak convergence arguments in our analysis; essentially the same issue appears in the papers [12, 28], which deal with more traditional mean field game models.

The paper is organized as follows. The next section presents the continuous-time model of bank run based on some of the ideas of [33] and Gossner’s lecture mentioned earlier. This is borrowed from the forthcoming book [11], and we present a streamlined version for the purpose of motivation. We use continuous time stochastic processes to model the value of the assets of the bank and the private signals of the depositors. Stylized facts from economic models of bank runs are captured in a set of assumptions about the costs and rewards to the depositors, and a mathematical problem of game of timing is articulated. Section 3 describes a general mathematical framework generalizing the set-up of the previous section. There, we provide all the required definitions and notation, and state the first main results of the paper, under complementarity assumptions. The following Sect. 4 specializes the setup further to models with continuous objective functions driven by Wiener processes. No proofs are given in these sections, only illustrations of how the abstract framework generalizes the bank run presented in Sect. 2, and how the results answer the questions raised therein. The remainder of the paper, from Sect. 5 on, is devoted to the proofs of the

results announced in Sects. 3 and 4. Section 6, in particular, develops the requisite material on filtration enlargements and randomized stopping times, some of which may be of independent interest. Two appendices provide proofs of technical results which we could not find in the printed literature.

2 A Model for Bank Runs

The nature of the balance sheet of a bank and the impact of the fire sales triggered by depositors runs and the possible failure of the bank are two important elements of the analysis of bank runs and their consequences, especially from a regulatory perspective. However, for the purpose of our mathematical analysis, we shall simplify their roles in order to focus on the optimal timing decisions of the investors.

Suppose the market value of the assets of a bank evolve over time according to some (real-valued) stochastic process $B = (B_t)_{t \geq 0}$, where the initial value $B_0 > 0$ of the bank assets is known to everyone, and in particular to the depositors. We assume that the assets generate a flow of dividends at rate $\bar{r}$ strictly greater than the risk free rate r . These dividends are not reinvested, so they are not included in B_t . The depositors are promised the same interest rate $\bar{r}$ on their deposits. The bank remains in business as long as $B_t > 0$.

Let n be the number of depositors. We shall eventually let $n \rightarrow \infty$ to derive a mean field game model. For this reason, we normalize the initial deposit of each investor to $D_0^i = 1/n$, so the aggregate initial deposit is 1. We introduce a (deterministic) function L (typically satisfying at least $L(0) = 0$ and $0 < L' < 1$), and we think of the value $L(B_t)$ as the *liquidation value* of the assets of the bank at time t . As L is deterministic, it is known to everyone.

Whenever an investor tries to withdraw his or her deposit, the bank taps a credit line at interest rate $\bar{r} > r$ to pay the running investor. At time t , the credit line limit is equal to the liquidation value $L(B_t)$ of the bank's assets. The model is set up this way to allow the bank to pay running investors without having to tinker with its investments.

The bank is said to be safe if all depositors can be paid in full, even in case of a run. The bank is said to have liquidity problems if the current market value of its assets is sufficient to pay depositors, but the liquidation value is not. Finally, it is said to be insolvent if the current market value of its assets is less than its obligation to depositors. We shall confirm below that, in the case of complete information about the value of the assets of the bank, depositors start to run as soon as the bank has liquidity problems, possibly long before the bank is insolvent.

Let T be a finite time horizon, for the sake of concreteness, but notice that the story to follow makes just as much sense with $T = \infty$ or even when T is an appropriately random time. At time T , the bank's assets mature and generate a single payoff B_T which can be used to pay the credit line and the depositors. Cash flows stop at time T . At that time,

- if $B_T \geq 1$, the bank is safe and everybody is paid in full;
- if $B_T < 1$, the bank cannot pay everybody in full, there is an *exogenous default*.

This is not the only way the bank can default. Indeed there is the possibility of an *endogenous default* at time $t < T$ if the aggregate amount of withdrawals by running

depositors exceeds $L(B_t)$. Let us denote by τ^i the time at which depositor i tries to withdraw his or her deposit, and by $\bar{\mu}^n$ the empirical distribution of these times, i.e.

$$\bar{\mu}^n = \frac{1}{n} \sum_{i=1}^n \delta_{\tau^i},$$

where we use the notation δ_x for the probability measure putting mass 1 on the singleton $\{x\}$. Notice that $\bar{\mu}^n[0, t)$ represents the proportion of depositors who tried to withdraw before time t , and that the time of endogenous default is given by¹

$$\tau^{endo} = \inf\{t \in (0, T); \bar{\mu}^n[0, t) > L(B_t)\},$$

with the convention that the infimum of the empty set is defined as T . For the sake of simplicity we assume that once a depositor runs, he cannot get back in the game, in other words, his decision is irreversible.

Depositor Strategic Behavior

We now explain the strategic behavior of the n depositors. We denote by $\mathbb{F}^i = (\mathcal{F}_t^i)_{t \geq 0}$ the information available to player $i \in \{1, \dots, n\}$. This is a filtration, $\mathcal{F}_t^i$ representing the information available to player i at time t . In the particular case which we discuss first, these filtrations are all identical. They are based on a perfect (though non-anticipative) observation of the signal $(B_t)_{0 \leq t \leq T}$. We call this situation *public monitoring*. In a more realistic form of the model, the filtration $\mathbb{F}^i$ is generated by the process $X^{i,n} = (X_t^{i,n})_{t \geq 0}$ and the process $(\bar{\mu}^n[0, t])_{t \geq 0}$, where $X_t^{i,n}$ is the *private signal* of depositor i , namely the value of the observation of B_t he or she can secure at time t . We shall assume that it is of the form

$$X_t^{i,n} = B_t + \sigma W_t^i$$

where $\sigma > 0$ and for $i \in \{1, \dots, n\}$, the processes $(W_t^i)_{t \geq 0}$ are independent identically distributed (i.i.d.) stochastic processes (also independent of B) representing idiosyncratic noise terms blurring the observations of the exact value B_t of the assets of the bank. When $\mathbb{F}^i$ is generated by $X^{i,n}$ and $(\bar{\mu}^n[0, t])_{t \geq 0}$, we talk about *private monitoring* of the asset value of the bank. However, for an even more realistic form of the model, we shall require that the filtration $\mathbb{F}^i$ is generated simply by $X^{i,n}$ and does not include the information provided by the process $(\bar{\mu}^n[0, t])_{t \geq 0}$, which incorporates the private signals of the other depositors. This model should be more challenging mathematically as the individual depositors will have to choose their withdrawal strategies in a distributed manner, using only the information contained in their private signals, ignoring the process $(\bar{\mu}^n[0, t])_{t \geq 0}$.

In any case, the filtrations $\mathbb{F}^i$ will be specified in each particular application and will play the following role: the time τ^i chosen by agent i is required to be a $\mathbb{F}^i$ -stopping time in order to be admissible.

¹ Here and throughout the text we write $\bar{\mu}^n[0, t)$ in place of the somewhat more precise $\bar{\mu}^n([0, t))$.

Given that all the other players $j \neq i$ have chosen their times τ^j to try to withdraw their deposits, the payoff $P^i(\tau^{-i}, \tau^i)$ to depositor i for trying to run on the bank at time τ^i can be written (recalling that $D_0^i = 1/n$) as

$$\begin{aligned} P^i(\tau^{-i}, \tau^i) &= D_0^i \wedge \left(L(B_{\tau^i}) - \bar{\mu}^n[0, \tau^i) \right)^+ \\ &= D_0^i \wedge \left(L(B_{\tau^i}) - \frac{1}{n} \sum_{k=1}^n \mathbf{1}_{[0, \tau^i)}(\tau^k) \right)^+, \end{aligned}$$

if $L(B_s) - \frac{1}{n} \sum_{k=1}^n \mathbf{1}_{[0, s)}(\tau^k) > 0$ for all $s < t$, and $P^i(\tau^{-i}, \tau^i) = 0$ otherwise. The problem of depositor i is then to choose for τ^i , the $\mathbb{F}^i$ -stopping time solving the maximization problem

$$J^i(\tau^{-i}) = \sup_{\tau^i} \mathbb{E} \left[e^{(\bar{r}-r)\tau^i} P^i(\tau^{-i}, \tau^i) \right]$$

which is an optimal stopping problem. Any solution τ^i of this maximization problem represents a best response of player i to the choices τ^{-i} of the other depositors. Finding a set of stopping times τ^i for $i = 1, \dots, n$ satisfactory to all the players simultaneously is essentially finding a fixed point to the search for best responses. This is achieved by finding a Nash equilibrium for this game.

Solution in the Case of Public Monitoring Through Perfect Observation

If we assume that $\sigma = 0$, in which case $\mathbb{F}^i = \mathbb{F}^B = (\mathcal{F}_t^B)_{t \geq 0}$, at time t each depositor knows the past up to time t of the asset value B_s for $s \leq t$, and if all the depositor decisions (to run or not to run) are based only on this information, then for each $t \in [0, T]$, $\bar{\mu}^n[0, t] \in \mathcal{F}_t^B$ since this information is known by all the depositors at time t .

Proposition 2.1 *In the case of public information, if we define the stopping time $\hat{\tau}$ by*

$$\hat{\tau} = T \wedge \inf \left\{ t > 0; L(B_t) \leq \frac{n-1}{n} \right\},$$

then a Nash equilibrium is when all the depositors decide to run at time $\hat{\tau}$.

So a bank run occurs as soon as the bank has liquidity problems, even if this is long before it is insolvent. Notice also that according to this proposition, all the depositors experience full recovery of their deposits, which is in flagrant contrast with typical bank runs in which most depositors usually experience significant losses.

Proof We first argue that we have indeed identified a Nash equilibrium. If all the other depositors but i choose the strategy given by the running time $\hat{\tau}$, we show that player i cannot do better than choosing to also run at time $\hat{\tau}$. If $L(B_{\hat{\tau}}) \leq (n-1)/n$, all the others depositors run immediately, and the only hope investor i has to get something out of his or her deposit is to run at time $\hat{\tau}$ as well. Now if $L(B_{\hat{\tau}}) > (n-1)/n$, no

depositor has a reason to run while $L(B_t) > (n-1)/n$ since by not running for a small time interval while $L(B_t)$ is still strictly greater than $(n-1)/n$, he or she can earn the superior interest $\bar{r} > r$ without facing any risk. This proves that every depositor using $\hat{\tau}$ as time to run is a Nash equilibrium. $\square$

The Mean Field Game Formulation

We now consider an asymptotic regime corresponding to a large number of depositors, sending $n \rightarrow \infty$, and we track the behavior of a representative depositor with initial deposit $D_0 > 0$. Although the payoffs P^i themselves decrease to zero, as $D_0^i = 1/n$, we are not terribly concerned with the asymptotic behavior of the *values* of the objective functions, so we may simply rescale P^i to nP^i in the n -player game without altering the set of equilibria. Indeed, the main quantity we wish to control in this asymptotic regime is the empirical distribution of the equilibrium stopping times, as this contains all of the information describing the timing of the bank run.

When n is large, the usual heuristics for mean field games suggest that, if the process B giving the asset value of the bank is not deterministic, $\bar{\mu}^n$ approaches a random measure μ . In particular, this limiting μ should depend on the time-evolution of B in the sense that $\mu[0, t]$ should be $\mathcal{F}_t^B$ -measurable for each $t \in [0, T]$. If such a probability measure μ is fixed, one defines the individual payoff $P^\mu(t, y)$ of a withdrawal attempt at time t when the value of the assets of the bank is y as:

$$P^\mu(t, y) = D_0 \wedge \left(L(y) - \mu[0, t] \right)^+,$$

as long as $L(B_s) - \mu[0, s]$ is positive for all $s < t$; otherwise the payoff is null. The optimal time for a representative depositor to claim his or her deposit back will be given by the stopping times solving the optimal stopping problem:

$$\sup_{0 \leq \theta \leq T} \mathbb{E}[e^{(\bar{r}-r)\theta} P^\mu(\theta, B_\theta)].$$

The above maximization is understood over all the $\mathbb{F}^X$ -stopping times θ where the filtration $\mathbb{F}^X = (\mathcal{F}_t^X)_{0 \leq t \leq T}$ is the filtration generated by the signal $X_t = B_t + \sigma W_t$ observed at time t by our generic investor. Here $(W_t)_{0 \leq t \leq T}$ is a process independent of $(B_t)_{0 \leq t \leq T}$ and sharing the same distribution as each of the W^i from before. If we can solve this optimal stopping problem for each (random) measure μ , we can define a map $\mu \rightarrow \text{Law}(\hat{\tau}|B)$ where $\hat{\tau}$ is an optimal stopping time, and the final step of the mean field game approach is to find a fixed point for this map. The following section formulates the mean field game more precisely and explains the connection with the n -player game.

3 General Mean Field Games of Timing: Main Results

A compact set of times $\mathbb{T} \subset [0, \infty]$ is fixed throughout, which we assume is either discrete or of the form $[0, T]$ for some $T \in [0, \infty]$. Fix two filtered probability spaces

$(\Omega^{\text{com}}, \mathcal{F}^{\text{com}}, \mathbb{F}^{\text{com}}, \mathbb{P}^{\text{com}})$ and $(\Omega^{\text{ind}}, \mathcal{F}^{\text{ind}}, \mathbb{F}^{\text{ind}}, \mathbb{P}^{\text{ind}})$, which will house a *common noise* and an *independent* (or *idiosyncratic*) *noise*, respectively. We are given also a filtration $\mathbb{F}^{\text{sig}} = (\mathcal{F}_t^{\text{sig}})_{t \in \mathbb{T}}$ on the product space $\Omega^{\text{com}} \times \Omega^{\text{ind}}$ with $\mathcal{F}_t^{\text{sig}} \subset \mathcal{F}_t^{\text{com}} \otimes \mathcal{F}_t^{\text{ind}}$ for every t . This filtration represents the *signal* or information available to an agent. Rather than observing the full filtration $\mathbb{F}^{\text{com}} \otimes \mathbb{F}^{\text{ind}}$ of the underlying noises, an agent sees only $\mathbb{F}^{\text{sig}}$. An *objective function* is given,

$$F : \Omega^{\text{com}} \times \Omega^{\text{ind}} \times \mathcal{P}(\mathbb{T}) \times \mathbb{T} \rightarrow \mathbb{R},$$

where $\mathcal{P}(\mathbb{T})$ denotes the set of Borel probability measures on $\mathbb{T}$. Here $F(\omega^0, \omega^1, m, t)$ represents the reward an agent achieves by stopping at time t , given the values (ω^0, ω^1) of the common and independent noises, and given the distribution m of other agents' stopping times.

With these ingredients, we will formulate both an n -player game and its continuum limit as $n \rightarrow \infty$. Assumption **A** below will clarify the precise assumptions (measurability, continuity, etc.) on F , and until then we will tacitly assume the expectations make sense.

Example 3.1 In the example presented in Sect. 2, we make the following identifications. Let $\mathbb{T} = [0, T]$, and let $(\Omega^{\text{com}}, \mathcal{F}^{\text{com}}, \mathbb{F}^{\text{com}}, \mathbb{P}^{\text{com}})$ and $(\Omega^{\text{ind}}, \mathcal{F}^{\text{ind}}, \mathbb{F}^{\text{ind}}, \mathbb{P}^{\text{ind}})$ both equal the Wiener space of continuous real-valued paths. That is, $\Omega^{\text{com}} = \Omega^{\text{ind}} = C([0, T])$ is equipped with the Borel σ -field, the Wiener measure, and the natural (augmented) filtration. The sub-filtration $\mathbb{F}^{\text{sig}}$ is the complete filtration generated by the process $(W_t^0 + \sigma W_t^1)_{t \in [0, T]}$, where W^0 and W^1 denote the projections from $\Omega^{\text{com}} \times \Omega^{\text{ind}}$ to Ω^{com} and to Ω^{ind} , respectively. The objective function (after a renormalization) is

$$F(\omega^0, \omega^1, m, t) = e^{(r-r)t} \left[1 \wedge \left(L(\omega_t^0) - m[0, t) \right)^+ \right],$$

if $L(\omega_s^0) - m[0, s) > 0$ for all $s < t$; otherwise, it is 0. Note that in this example the independent noise ω^1 does not appear in the payoff, and its only role is in specifying the information structure $\mathbb{F}^{\text{sig}}$.

3.1 The n -Player Game

The n -player game for $n \geq 1$ is defined on the product space

$$(\overline{\Omega}, \overline{\mathcal{F}}, \overline{\mathbb{F}}, \overline{\mathbb{P}}) := (\Omega^{\text{com}}, \mathcal{F}^{\text{com}}, \mathbb{F}^{\text{com}}, \mathbb{P}^{\text{com}}) \otimes \bigotimes_{k=1}^{\infty} (\Omega^{\text{ind}}, \mathcal{F}^{\text{ind}}, \mathbb{F}^{\text{ind}}, \mathbb{P}^{\text{ind}}).$$

A typical element of $\overline{\Omega}$ is denoted $\omega = (\omega^0, \omega^1, \dots)$, with $\omega^0 \in \Omega^{\text{com}}$ and $\omega^i \in \Omega^{\text{ind}}$ for $i \geq 1$. We call ω^0 the *common noise* and ω^i the *idiosyncratic noise* of agent i . Define the projections

$$W^i(\omega^0, \omega^1, \dots) = \omega^i,$$

for $i = 0, 1, \dots$. Finally, for $i \geq 1$, define the filtration $\mathbb{F}^i = (\mathcal{F}_t^i)_{t \in \mathbb{T}}$ of the i th agent by

$$\mathcal{F}_t^i := (W^0, W^i)^{-1}(\mathcal{F}_t^{\text{sig}}) := \sigma \left\{ \left\{ (W^0, W^i) \in C \right\} : C \in \mathcal{F}_t^{\text{sig}} \right\}.$$

Define the empirical measure map $\bar{\mu}^n : \mathbb{T}^n \rightarrow \mathcal{P}(\mathbb{T})$ by

$$\bar{\mu}^n(t_1, \dots, t_n) = \frac{1}{n} \sum_{k=1}^n \delta_{t_k}. \quad (3.1)$$

We will make use of the following common notation: given $\mathbf{e} = (e_1, \dots, e_n) \in E^n$ for some set E , define

$$\begin{aligned} \mathbf{e}^{-i} &= (e_1, \dots, e_{i-1}, e_{i+1}, \dots, e_n), \quad \text{and} \quad (\mathbf{e}^{-i}, x) \\ &= (e_1, \dots, e_{i-1}, x, e_{i+1}, \dots, e_n), \end{aligned}$$

for $x \in E$ and $i = 1, \dots, n$. To minimize the number of parentheses, we abuse notation somewhat by writing $\bar{\mu}^n(\mathbf{t}^{-i}, s)$ in lieu of $\bar{\mu}^n((\mathbf{t}^{-i}, s))$, when $\mathbf{t} \in \mathbb{T}^n$ and $s \in \mathbb{T}$. For $\epsilon \geq 0$, we say that $\boldsymbol{\tau} = (\tau_1, \dots, \tau_n)$ is an ϵ -Nash equilibrium if τ_i is an $\mathbb{F}^i$ -stopping time (defined on $\bar{\Omega}$) and

$$\mathbb{E} \left[F(W^0, W^i, \bar{\mu}^n(\boldsymbol{\tau}), \tau_i) \right] \geq \mathbb{E} \left[F(W^0, W^i, \bar{\mu}^n(\boldsymbol{\tau}^{-i}, \sigma), \sigma) \right] - \epsilon,$$

for every alternative $\mathbb{F}^i$ -stopping time σ , for each $i = 1, \dots, n$.

3.2 The Mean Field Game

We next define the infinite-agent counterpart of the above game, called the *mean field game*, which is formulated on the product space $\Omega^{\text{com}} \times \Omega^{\text{ind}}$. We write $\mathbb{E}$ for expectation under the product measure $\mathbb{P}^{\text{com}} \times \mathbb{P}^{\text{ind}}$, and we write W^0 and W^1 for the projections onto Ω^{com} and Ω^{ind} , respectively.

Definition 3.2 A strong mean field equilibrium (MFE) is a $\mathbb{F}^{\text{sig}}$ -stopping time τ^* on $\Omega^{\text{com}} \times \Omega^{\text{ind}}$ satisfying

$$\mathbb{E} \left[F(W^0, W^1, \mu, \tau^*) \right] \geq \mathbb{E} \left[F(W^0, W^1, \mu, \tau) \right],$$

for every alternative $\mathbb{F}^{\text{sig}}$ -stopping time τ , where

$$\mu = \mathbb{P}^{\text{com}} \times \mathbb{P}^{\text{ind}}[\tau^* \in \cdot \mid W^0] \quad (3.2)$$

is the regular conditional law of τ^* given W^0 .

We say *strong MFE* here because later, in Definition 4.2, we will later introduce a notion of *weak MFE*. One justification of this strong equilibrium concept is the following theorem, which explains how to use a strong MFE to construct approximate Nash equilibria for the n -player games. First, some assumptions are needed. In the following, consider the topology $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ generated by the set $B(\mathbb{T})$ of bounded measurable functions of $\mathbb{T}$; that is, $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ is the coarsest topology on $\mathcal{P}(\mathbb{T})$ such that the map $m \mapsto \int_{\mathbb{T}} \varphi dm$ is continuous for every $\varphi \in B(\mathbb{T})$. Define the total variation of a signed measure ν on $\mathbb{T}$ by

$$\|\nu\|_{TV} = \sup \left\{ \int_{\mathbb{T}} f d\nu : f \in B(\mathbb{T}), \sup_{t \in \mathbb{T}} |f(t)| \leq 1 \right\}.$$

In the following, $\Omega^{\text{com}} \times \Omega^{\text{ind}}$ is always equipped with the probability measure $\mathbb{P}^{\text{com}} \times \mathbb{P}^{\text{ind}}$.

Assumption A

- (A.1) F is jointly measurable, with $\mathcal{P}(\mathbb{T})$ equipped with the σ -field generated by the maps $m \mapsto m(C)$, where $C \subset \mathbb{T}$ is a Borel set.²
- (A.2) For almost every (ω^0, ω^1) , the map $m \mapsto F(\omega^0, \omega^1, m, t)$ is $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ -continuous, uniformly in t . That is, for each $m_0 \in \mathcal{P}(\mathbb{T})$, the map

$$m \mapsto \sup_{t \in \mathbb{T}} \left| F(\omega^0, \omega^1, m, t) - F(\omega^0, \omega^1, m_0, t) \right|$$

is $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ -continuous at m_0 .

- (A.3) It holds that

$$\mathbb{E} \left[\sup_{m \in \mathcal{P}(\mathbb{T})} \sup_{t \in \mathbb{T}} \left| F(W^0, W^1, m, t) \right| \right] < \infty. \quad (3.3)$$

Assumption (A.2) may appear difficult to verify. On the contrary, there are two broad classes of examples it covers. First, because $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ is finer than the topology of weak convergence, replacing $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ -continuity with weak continuity (i.e., continuity with respect to the topology of weak convergence) is enough. Moreover, because $\mathbb{T}$ is compact, joint continuity of $F(\omega^0, \omega^1, m, t)$ in (m, t) implies (A.2). The second class of examples, and indeed the one that motivates the use of the topology $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$, consists of functions F of the form $F(\omega^0, \omega^1, m, t) = G(\omega^0, \omega^1, m[0, t], t)$, where $G : \Omega^{\text{com}} \times \Omega^{\text{ind}} \times [0, 1] \times \mathbb{T} \rightarrow \mathbb{R}$ is measurable. If $G = G(\omega^0, \omega^1, u, t)$ is continuous in u , uniformly in t , for each fixed (ω^0, ω^1) , then F satisfies (A.2). This follows from a simple lemma, proven in Sect. 5.

² This σ -field agrees with the Borel σ -field generated by the topology of weak convergence on $\mathcal{P}(\mathbb{T})$.

Lemma 3.3 For each $m_0 \in \mathcal{P}(\mathbb{T})$, the map

$$m \mapsto \sup_{t \in \mathbb{T}} |m[0, t] - m_0[0, t]|$$

is $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ -continuous at m_0 .

The following main result illustrates how a mean field equilibrium can be used to construct near-equilibria for the n -player games. Its proof is given in Sect. 5.

Theorem 3.4 Suppose assumption A holds. Suppose τ^* is a mean field equilibrium, and let μ be as in (3.2). For each k define an $\mathbb{F}^k$ -stopping time on $\bar{\Omega}$ by

$$\tau^k(\omega^0, \omega^1, \dots, \omega^n) = \tau^*(\omega^0, \omega^k).$$

Then there exist $\epsilon_n \geq 0$ with $\epsilon_n \rightarrow 0$ such that $\tau^n = (\tau^1, \dots, \tau^n)$ is an ϵ_n -Nash equilibrium for each n , and moreover

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[F(W^0, W^k, \bar{\mu}^n(\tau^n), \tau^k) \right] = \mathbb{E} \left[F(W^0, W^1, \mu, \tau^*) \right], \text{ for each } k.$$

3.3 Strategic Complementarities and Existence of MFEs

An existence result for strong MFE is available, even for discontinuous F , as long as a suitable *complementarity* property holds, as was mentioned in the introduction. This section draws heavily on ideas from literature on games with strategic complementarities [30, 37], which is abundant with existence proofs based more on monotonicity than continuity. In the following, let us say that a $\mathcal{P}(\mathbb{T})$ -valued random variable μ on $\Omega^{\text{com}} \times \Omega^{\text{ind}}$ is an $\mathbb{F}^{\text{com}}$ -adapted random measure if $\mu[0, t]$ is $\mathcal{F}_t^{\text{com}}$ -measurable for every $t \in \mathbb{T}$.

Assumption B

(B.1) $\mathbb{F}^{\text{sig}}$ is right-continuous.

(B.2) For every pair of $\mathbb{F}^{\text{com}}$ -adapted random measures $\mu, \tilde{\mu}$ satisfying $\mu[0, t] \geq \tilde{\mu}[0, t]$ for all $t \in \mathbb{T}$ a.s., the process $(M_t)_{t \in \mathbb{T}}$ defined by

$$M_t = F(W^0, W^1, \tilde{\mu}, t) - F(W^0, W^1, \mu, t)$$

is a submartingale.

(B.3) For each $m \in \mathcal{P}(\mathbb{T})$, $t \mapsto F(W^0, W^1, m, t)$ is upper semicontinuous, almost surely.

(B.4) Conditions (A.1) and (A.3) hold.

If $\mu \leq \tilde{\mu}$ in the sense of stochastic order (i.e., if $\tilde{\mu}[0, t] \leq \mu[0, t]$ a.s. for each $t \in [0, T]$), and if $\tau \leq \tilde{\tau}$ are stopping times, taking expectations in the submartingale property of M_t in assumption (B.2) yields

$$\begin{aligned} & \mathbb{E}[F(W^0, W^1, \tilde{\mu}, \tilde{\tau})] - \mathbb{E}[F(W^0, W^1, \tilde{\mu}, \tau)] \\ & \geq \mathbb{E}[F(W^0, W^1, \mu, \tilde{\tau})] - \mathbb{E}[F(W^0, W^1, \mu, \tau)], \end{aligned} \quad (3.4)$$

property which is usually called *increasing differences*. Intuitively, assumption (B.2) requires that for “larger” μ the function F increases more rapidly in expectation with t than it does for smaller μ . These hypotheses introduce strategic complementarities in the game and recast the game of timing model as a supermodular game. They are natural in the context of bank run models, in which the measure μ captures how early people run to the bank. Indeed, if $\tilde{\mu} \geq \mu$ in stochastic order, then under μ more people have run to the bank earlier. Under μ , the reward an agent gains by waiting from τ to $\tilde{\tau} > \tau$ should not exceed the same reward under $\tilde{\mu}$. In other words, if people tend to run to the bank earlier, the “cost of waiting” for an investor should be greater.

While assumption **B** is all that is needed for existence, the following stronger assumption will enable a better understanding of the structure of equilibria.

Assumption C

(C.1) $F(W^0, W^1, m, t)$ is almost surely jointly continuous in (m, t) , when $\mathcal{P}(\mathbb{T})$ is endowed with the topology of weak convergence.

(C.2) Condition (A.3) holds.

Theorem 3.5 *If assumption B holds, then there exists a strong MFE. If both assumptions B and C hold, then there exist strong MFEs τ^* and θ^* such that for any strong MFE τ we have $\theta^* \leq \tau \leq \tau^*$ a.s.*

Some examples of assumption **B** are as follows. First, suppose that for every $t \leq t'$, every (ω^0, ω^1) , and every $m, m' \in \mathcal{P}(\mathbb{T})$ satisfying $m \leq m'$ in stochastic order (meaning $m[0, s] \geq m'[0, s]$ for every s), we have

$$F(\omega^0, \omega^1, m', t') - F(\omega^0, \omega^1, m', t) \geq F(\omega^0, \omega^1, m, t') - F(\omega^0, \omega^1, m, t).$$

Then the submartingale part of assumption **B** holds trivially, as the process M is nondecreasing. The following proposition and remark show how to verify assumption **B** for a large class of examples based on diffusion processes.

Proposition 3.6 *Suppose $\mathbb{T} = [0, T]$, and assume the space $\Omega^{\text{com}} \times \Omega^{\text{ind}}$ supports a continuous Itô diffusion $X = (X_t)_{t \in [0, T]}$ with infinitesimal generator $\mathcal{L}$ defined on all smooth functions φ of compact support by*

$$\mathcal{L}f(x) = b(x) \cdot \nabla f(x) + \frac{1}{2} \text{Tr}[a(x) \nabla^2 f(x)],$$

where b and a are measurable functions with values in $\mathbb{R}^d$ and the set of positive semidefinite $d \times d$ matrices, respectively. Assume F is of the form $F(\omega^0, \omega^1, m, t) = f(X_t(\omega^0, \omega^1), \varphi * m(t), t)$, where f is bounded and

$$\varphi * m(t) = \int_{[0, T]} \varphi(t - s) m(ds).$$

Moreover, assume $f : \mathbb{R}^d \times \mathbb{R} \times [0, T] \ni (x, y, t) \mapsto f(x, y, t) \in \mathbb{R}$ has two bounded continuous derivatives in x and one in both y and t , and $\varphi : [-T, T] \rightarrow \mathbb{R}$ is continuous. Suppose one of the following holds:

- (i) φ is nondecreasing and convex, $\partial_y f \geq 0$, and also $\mathcal{L}_x f + \partial_t f$ and $\partial_y f$ are nondecreasing in y for each fixed (x, t) , where $\mathcal{L}_x$ denotes the action of $\mathcal{L}$ on the x variable.
- (ii) φ is nonincreasing and convex, $\partial_y f \leq 0$, and also $\mathcal{L}_x f + \partial_t f$ and $\partial_y f$ are nonincreasing in y for each fixed (x, t) .

Assume finally that

$$\mathbb{E} \int_0^T \sup_{y \in \mathbb{R}} |a(X_t) \nabla_x f(X_t, y, t)|^2 dt < \infty. \quad (3.5)$$

Then assumptions (B.2-3) and (C.1) hold.

Proof The only nontrivial claim is that the submartingale property (B.2) holds. To check this, fix two $\mathbb{F}^{\text{com}}$ -adapted random measures μ and $\tilde{\mu}$ satisfying $\mu[0, t] \geq \tilde{\mu}[0, t]$ a.s. for every t . By Itô's formula,

$$\begin{aligned} df(X_t, \varphi * \mu(t), t) = & \{ \mathcal{L}_x f(X_t, \varphi * \mu(t), t) + \partial_t f(X_t, \varphi * \mu(t), t) \\ & + \partial_y f(X_t, \varphi * \mu(t), t) \varphi' * \mu(t) \} dt \\ & + \nabla_x f(X_t, \varphi * \mu(t), t) \cdot a(X_t) dB_t, \end{aligned}$$

where $B = (B_t)_{t \in [0, T]}$ is a standard Brownian motion (defined perhaps on an extension of the probability space). The assumption (3.5) implies that the dB_t term is a martingale. To show that $f(B_t, \varphi * \mu(t), t) - f(B_t, \varphi * \tilde{\mu}(t), t)$ is a submartingale, it suffices to check that its dt term is always nonnegative. If $\mu \leq \tilde{\mu}$ are as in assumption B, then the dt term of $df(B_t, \varphi * \mu(t), t) - df(B_t, \varphi * \tilde{\mu}(t), t)$ is precisely

$$\begin{aligned} & \mathcal{L}_x f(B_t, \varphi * \tilde{\mu}(t), t) - \mathcal{L}_x f(B_t, \varphi * \mu(t), t) \\ & + \partial_t f(B_t, \varphi * \tilde{\mu}(t), t) - \partial_t f(B_t, \varphi * \mu(t), t) \\ & + \partial_y f(B_t, \varphi * \tilde{\mu}(t), t) \varphi' * \tilde{\mu}(t) - \partial_y f(B_t, \varphi * \mu(t), t) \varphi' * \mu(t). \end{aligned}$$

Now note that if $m \leq \tilde{m}$ in stochastic order then $\int g dm \leq \int g d\tilde{m}$ for every nondecreasing function g , and in particular if φ is nondecreasing (resp. nonincreasing) and convex then $\varphi * \tilde{m} \geq \varphi * m$ (resp. $\leq$) and $\varphi' * \tilde{m} \geq \varphi' * m$ pointwise. With this in mind, it is straightforward to check that either set of assumptions ensures that the above quantity is nonnegative. $\square$

The assumption (3.5) is not very restrictive; it holds as soon as $\nabla_x f$ and a are bounded, or more generally under linear growth assumptions and suitable integrability for the initial state X_0 . The conditions (i–ii) are more restrictive, and the following simple result illustrates more broadly the limitations of assumption (B.2) in handling a very natural form of mean field interaction. In particular, Proposition 3.7 suggests that our bank run model cannot satisfy assumption (B.2) because of the dependence of $F(\omega^0, \omega^1, m, t)$ on $m[0, t]$.

Proposition 3.7 *Suppose $\mathbb{T}$ is continuous, and suppose $F(\omega^0, \omega^1, m, t) = G(m[0, t])$ for some continuous $G : [0, 1] \rightarrow \mathbb{R}$ which we assume is differentiable on $(0, 1)$. If F satisfies assumption (B.2), then G is constant.*

Proof For $m, \tilde{m} \in \mathcal{P}([0, T])$ with $m \leq \tilde{m}$, assumption (B.2) implies that the deterministic process $G(\tilde{m}[0, t]) - G(m[0, t])$ is a submartingale, which means simply that it is nondecreasing. In other words, for $0 \leq s \leq t \leq T$,

$$G(\tilde{m}[0, t]) - G(\tilde{m}[0, s]) \geq G(m[0, t]) - G(m[0, s]).$$

Dividing by $t - s$ and taking limits, we find

$$G'(F_2(t))f_2(t) \geq G'(F_1(t))f_1(t),$$

assuming $F_1(t) = m[0, t]$ and $F_2(t) = \tilde{m}[0, t]$ have derivatives f_1 and f_2 . The point is that stochastic dominance is not sensitive to changes in density. Given $u \in (0, 1)$, there exist $m \leq \tilde{m}$ and $t \in [0, T]$ such that $F_2(t) = u$ while $f_1(t) = 0$ and $f_2(t) = 1$, which implies $G'(u) \geq 0$. On the other hand, given $u \in (0, 1)$, there exist $m \leq \tilde{m}$ and $t \in [0, T]$ such that $F_1(t) = u$ while $f_1(t) = 1$ and $f_2(t) = 0$, which implies $G'(u) \leq 0$. Thus $G' \equiv 0$ on $(0, 1)$. $\square$

3.4 An Example

The recent model of Nutz [32], or at least many specializations thereof, can be shown to satisfy our assumption A. The explicit computations of equilibria in [32] can be used in tandem with Theorem 3.4 to construct n -player approximate equilibria.

We describe only a simple case of this model, from [32, Sect. 5.1]. Let $\mathbb{T} = [0, \infty]$, and suppose $\Omega^{\text{com}} = \Omega^{\text{ind}} = D_{\uparrow}$ is the space of nondecreasing right-continuous real-valued functions on $[0, \infty)$. Note that for any $f \in D_{\uparrow}$ the limit $f(\infty) = \lim_{t \rightarrow \infty} f(t)$ exists in $\mathbb{R} \cup \{\infty\}$. Constants $c > 0$ and $r > 0$ are given, and the objective function is

$$F(\omega^0, \omega^1, m, t) = \exp \left(\int_0^t \left(r - \omega^0(s) - \omega^1(s) - cm[0, s] \right) ds \right).$$

The process $\gamma_t(\omega^0, \omega^1) = \omega^0(t) - \omega^1(t) - cm[0, t]$ can be interpreted as the agent's perception of the rate of bank failure. This perceived rate changes over time, depending on a common factor ω^0 and an independent factor ω^1 , as well as the fraction of agents who have already run to the bank. In fact, this is not the primitive form of the objective function given in [32] but is instead derived in Lemma 2.1 therein (more precisely, Eq. (2.3)).

It is straightforward to check that assumptions (A.1) holds for this example, and we may use Lemma 3.3 to check that (A.2) holds as well. Assumption (A.3) holds as long as

$$\mathbb{E} \left[\exp \left(\sup_{t \geq 0} \int_0^t \left(r - W^0(s) - W^1(s) \right) ds \right) \right] < \infty.$$

On the other hand, it appears that assumption (B.2) fails for this class of models in most cases. However, the arguments of [32] lead to explicit computations of MFE when agents have access to enough information, namely when W^0 and $W^0 + W^1$ are both adapted to $\mathbb{F}^{\text{sig}}$. Theorem 3.4 can then be used to construct explicit n -player approximate equilibria.

4 Beyond Complementarities: Weak Equilibria

This section explains how to move past the restrictive assumptions of complementarities by deriving an existence result and a limit theorem under the modest continuity assumptions on the objective function. Our time set is now a finite interval $\mathbb{T} = [0, T]$, $T > 0$. Let $\mathcal{C} = C([0, T])$ denote the space of continuous real-valued functions on $[0, T]$, endowed with the supremum norm. For a Polish space E , we always write $\mathcal{P}(E)$ for the space of Borel probability measures on E , endowed with the topology of weak convergence.

For the rest of this section we specify

$$\Omega^{\text{com}} = \Omega^{\text{ind}} = \mathcal{C},$$

The common noise and independent noise will now both be one-dimensional standard Brownian motions, for the sake of simplicity. This could be generalized in various directions, most obviously by making these Brownian motions multi-dimensional, and this would not alter the analysis. Let us write $\mathcal{W}$ for the Wiener measure on $\mathcal{C}$, and specialize the setup of Sect. 3 by setting

$$\mathbb{P}^{\text{com}} = \mathbb{P}^{\text{ind}} = \mathcal{W}.$$

Write $B = (B_t)_{t \in [0, T]}$ and $W = (W_t)_{t \in [0, T]}$ for the canonical processes on $\mathcal{C}^2$, and let $\mathbb{F}^B = (\mathcal{F}_t^B)_{t \in [0, T]}$ and $\mathbb{F}^W = (\mathcal{F}_t^W)_{t \in [0, T]}$ denote their natural (raw) filtrations. The objective function is now a function $F : \mathcal{C}^2 \times \mathcal{P}([0, T]) \times [0, T] \rightarrow \mathbb{R}$. Note that the full information version of the bank run model of Sect. 2 fits into this specialized setup; see Example 3.1.

The equilibrium concept for the n -player game is as in Sect. 3.1, but now with full information: Given independent Wiener processes B and $(W^i)_{i=1}^n$, agent i chooses a random time τ^i , which is required to be a stopping time relative to the full filtration $\mathbb{F}^{B, W^1, \dots, W^n}$ generated by $(B, W^1, \dots, W^n)$, but we will not spell out the details. Recall that for $\epsilon \geq 0$ we say that $\tau = (\tau_1, \dots, \tau_n)$ is an ϵ -Nash equilibrium if τ_i is an $\mathbb{F}^{B, W^1, \dots, W^n}$ -stopping time (with values in $[0, T]$) and if

$$\mathbb{E} \left[F \left(B, W^i, \bar{\mu}^n(\tau), \tau_i \right) \right] \geq \mathbb{E} \left[F \left(B, W^i, \bar{\mu}^n(\tau^{-i}, \sigma), \sigma \right) \right] - \epsilon,$$

for every alternative $\mathbb{F}^{B, W^1, \dots, W^n}$ -stopping time σ , for each $i = 1, \dots, n$. Unfortunately, our proof techniques seem to be restricted to this full information case; an

earlier version of this paper contained partial-information analogs of the following results, but there was a flaw in the proof.

We are interested in describing the limiting behavior of Nash equilibria, as $n \rightarrow \infty$, in addition to the converse construction of Theorem 4.5. To this end, we introduce notions of strong and weak equilibria in analogy with strong and weak solutions of stochastic differential equations. The strong equilibrium is exactly as in Definition 3.2, but with full information:

Definition 4.1 A *strong mean field equilibrium (MFE)* is a $\mathbb{F}^{B,W}$ -stopping time τ^* defined on $\mathcal{C}^2$, equipped with the Wiener measure $\mathcal{W}^2$, satisfying

$$\mathbb{E}[F(B, W, \mu, \tau^*)] \geq \mathbb{E}[F(B, W, \mu, \tau)],$$

for every $\mathbb{F}^{B,W}$ -stopping time τ , where $\mu = \mathcal{W}^2[\tau^* \in \cdot | B]$ is the conditional law of τ^* given B .

The definition of a weak MFE requires care. Because we will work heavily with weak limits, we must prepare for some loss of measurability, in light of the following basic fact of weak convergence: if (Z, Y_n) are random variables converging weakly to (Z, Y) , and if Y_n is Z -measurable for each n , then there is absolutely no reason to expect that Y is Z -measurable in the limit, despite the fact that Z does not depend on n . For this reason, we define a notion of weak MFE in which μ is not required to be B -measurable, and τ may be a *randomized stopping time*, in a sense made precise below. In analogy with the definition of weak solutions for stochastic differential equations, we base the definition of weak solution on the properties of the joint distribution of (B, W, μ, τ) . In fact, for reasons which will become clear later, it is convenient to include more information by considering not only the conditional law of τ but rather the joint conditional law of (W, τ) . Hence, we work with the canonical space

$$\Omega := \mathcal{C}^2 \times \mathcal{P}(\mathcal{C} \times [0, T]) \times [0, T], \quad (4.1)$$

and let (B, W, μ, τ) denote the canonical process given by the natural projections, $(B, W) : \Omega \rightarrow \mathcal{C}^2$, $\mu : \Omega \rightarrow \mathcal{P}(\mathcal{C} \times [0, T])$, and $\tau : \Omega \rightarrow [0, T]$. Because we will work with a number of canonical filtrations on this space, we introduce the following notation which we shall use systematically in the sequel. The continuous processes B and W each generate filtrations $\mathbb{F}^B = (\mathcal{F}_t^B)_{t \in [0, T]}$ and $\mathbb{F}^W$, respectively, defined in the natural way. The random time τ generates the raw filtration

$$\mathcal{F}_t^\tau = \sigma\{\tau \wedge t\}.$$

The filtration generated by multiple processes is denoted, for instance, by $\mathbb{F}^{B,W} := \mathbb{F}^B \vee \mathbb{F}^W$, or $\mathcal{F}_t^{B,W} = \sigma(\mathcal{F}_t^B \cup \mathcal{F}_t^W)$. We use the same notation $\mathbb{F}^{W,\tau}$ not only for the filtration $\mathbb{F}^W \vee \mathbb{F}^\tau$ defined on Ω , but also for the filtration generated on $\mathcal{C} \times [0, T]$, and this should not cause any confusion. With this identification, the filtration $\mathbb{F}^\mu = (\mathcal{F}_t^\mu)_{t \in [0, T]}$ on Ω (or on $\mathcal{P}(\mathcal{C} \times [0, T])$) is defined by

$$\mathcal{F}_t^\mu = \sigma\{\mu(C) : C \in \mathcal{F}_t^{W,\tau}\}.$$

Equivalently, if π_t is defined on $\mathcal{C} \times [0, T]$ by $\pi_t(w, s) = (w_{\cdot \wedge t}, s \wedge t)$, then $\mathcal{F}_t^\mu = \sigma\{\mu \circ \pi_t^{-1}\}$. We also write

$$\mu^W = \mu(\cdot \times [0, T]), \quad \mu^\tau = \mu(\mathcal{C} \times \cdot),$$

for the two marginals of μ , which take values in $\mathcal{P}(\mathcal{C})$ and $\mathcal{P}([0, T])$, respectively. Given a filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$, we write $\mathbb{F}_+$ for the right-continuous filtration $(\mathcal{F}_{t+})_{t \in [0, T]}$, where as usual $\mathcal{F}_{t+} := \cap_{s>t} \mathcal{F}_s$ for $t \in [0, T)$ and $\mathcal{F}_{T+} = \mathcal{F}_T$. Note that the right-filtration $\mathbb{F}_+$ is the smallest filtration for which τ is a stopping time, and for this reason the appearance of right-continuous filtrations in the following definition quite natural:

Definition 4.2 A *weak mean field equilibrium (MFE)* is a probability measure P on Ω such that:

- (1) (B, W) is a Wiener process with respect to the full filtration $\mathbb{F}_+^{B, W, \mu, \tau}$.
- (2) (B, μ) is independent of W .
- (3) τ is compatible with (B, W, μ) , in the sense that $\mathcal{F}_{t+}^\tau$ is conditionally independent of $\mathcal{F}_T^{B, W, \mu}$ given $\mathcal{F}_{t+}^{B, W, \mu}$, for every $t \in [0, T]$.
- (4) The optimality condition holds:

$$\mathbb{E}^P[F(B, W, \mu^\tau, \tau)] = \sup_{P'} \mathbb{E}^{P'}[F(B, W, \mu^\tau, \tau)],$$

where the supremum is over all $P' \in \mathcal{P}(\Omega)$ satisfying (1-3) as well as $P' \circ (B, W, \mu)^{-1} = P \circ (B, W, \mu)^{-1}$.

- (5) The weak fixed point condition holds: $\mu = P((W, \tau) \in \cdot \mid B, \mu)$.

The following result is the first justification for the above definition, and after stating it we will elaborate further on the intuitive meaning of a weak MFE. Recall that $\mathcal{W}$ denotes Wiener measure, and write $\mathcal{W}^2 = \mathcal{W} \times \mathcal{W}$ for the product measure on $\mathcal{C}^2$.

Proposition 4.3 Assume that F is bounded and jointly measurable and that $t \mapsto F(b, w, m, t)$ is continuous, for every m and $\mathcal{W}^2$ -almost every (b, w) . Suppose τ^* is a strong MFE, and define $\mu = \mathcal{W}^2(\tau^* \in \cdot \mid B)$. Then the measure

$$P = \mathcal{W}^2 \circ (B, W, \mu, \tau^*)^{-1} \quad (4.2)$$

is a weak MFE.

The proof of Proposition 4.3 is in Sect. 6. With some abuse of terminology, we may refer to the measure P itself, defined in (4.2), as a strong MFE. We may define also some intermediate notions of MFE. It may happen that τ is a.s. (B, W, μ) -measurable under P , in which case we say P is a *weak MFE with strong stopping time*.³ In contrast,

³ To say that τ is a.s. (B, W, μ) -measurable under P means that τ is measurable with respect to the P -completion of $\sigma(B, W, \mu)$. Equivalently, there exists a measurable map $\tilde{\tau} : \mathcal{C}^2 \times \mathcal{P}(\mathcal{C} \times [0, T]) \rightarrow [0, T]$ such that $P(\tau = \tilde{\tau}(B, W, \mu)) = 1$.

we may refer to a weak MFE more verbosely as a *weak MFE with weak stopping time*, to emphasize the failure of τ to be (B, W, μ) -measurable. Likewise, we say that a weak MFE P is a *strong MFE with weak stopping time* if μ is P -a.s. B -measurable. A *strong MFE with strong stopping time*, naturally, requires both of these measurability conditions, and according to Proposition 4.3 this reduces to what we have already called a *strong MFE*.

The “compatibility” condition (3) of Definition 4.2 is somewhat unusual. As mentioned above, we cannot expect τ to be (B, W, μ) -measurable after taking weak limits, but conditions (3) captures an important structure we do retain, as does the requirement in (1) that (B, W) remain Wiener processes with respect to the larger filtration. Similar compatibility conditions were identified in the stochastic differential mean field games in [12, 28] (see also [11]), and indeed these notions of compatibility all fall under the same umbrella, which we clarify somewhat in Sect. 6. Intuitively, our representative agent is allowed to randomize her stopping time externally to the signal (B, W, μ) , as long as at each time t this randomization is conditionally independent of all future information given the history of the signal. Mathematically, the reason compatibility arises is the following, stated informally here and made precise in Theorem 6.4: given τ satisfying (3), there exists a sequence of $\mathbb{F}^{B, W, \mu}$ -stopping times τ_k such that $(B, W, \mu, \tau_k) \Rightarrow (B, W, \mu, \tau)$, where $\Rightarrow$ denotes convergence in law. In this sense, the set of compatible stopping times is the closure of the set of *bona fide* stopping times.

Continuous Objective Functions

We are nearly ready to state the main results of this section, but first we need some assumptions:

Assumption D The function F is bounded and jointly measurable, and $\mathcal{P}([0, T]) \times [0, T] \ni (m, t) \mapsto F(b, w, m, t)$ is continuous for $\mathcal{W}^2$ -almost every $(b, w) \in \mathcal{C}^2$, when $\mathcal{P}([0, T])$ is equipped with the topology of weak convergence.

The boundedness assumption is for convenience only, and this could easily be relaxed at the cost of some careful growth or integrability assumptions. The continuity assumption is important for our weak convergence methods, but unfortunately it can be restrictive. For instance, our bank run model in the introduction involved the discontinuous function $\mathcal{P}([0, T]) \times [0, T] \ni (m, t) \mapsto m[0, t]$. A close approximation of the bank run model could be accounted for nonetheless by replacing $m[0, t]$ by $\phi * m(t) = \int_{[0, T]} \phi(t - s)m(ds)$, where $\phi : [-T, T] \rightarrow \mathbb{R}$ is continuous and in some sense “close to” the step function $1_{[0, T]}$.

The first result is a limit theorem, stating that n -player equilibria converge to weak MFE. Recall the notation of the n -player game in Sect. 3.1. For each n and each $t^1, \dots, t^n \in [0, T]$ we define the random joint empirical measure (a measure on $\mathcal{C} \times [0, T]$)

$$\widehat{\mu}^n(t^1, \dots, t^n) = \frac{1}{n} \sum_{i=1}^n \delta_{(W^i, t^i)}. \quad (4.3)$$

Theorem 4.4 Suppose assumption **D** holds. Let $\epsilon_n \geq 0$ with $\epsilon_n \rightarrow 0$, and suppose $\tau^n = (\tau_1^n, \dots, \tau_n^n)$ is an ϵ_n -Nash equilibrium for the n -player game for each n . Define

$$P_n = \frac{1}{n} \sum_{i=1}^n \mathbb{P} \circ \left(B, W^i, \widehat{\mu}^n(\tau^n), \tau_i^n \right)^{-1}.$$

Then $(P_n)_{n=1}^\infty$ is tight, and every weak limit is a weak MFE.

The measure P_n appearing in Theorem 4.4 is quite a natural object to study, if interpreted the right way. We may write $P_n = \mathbb{P} \circ (B, W^U, \widehat{\mu}^n(\tau^n), \tau_U^n)^{-1}$, where U is a random variable drawn uniformly from $\{1, \dots, n\}$, independent of $(B, W^i)_{i=1}^\infty$. Think of this as a *randomly selected representative agent*. As τ_i^n may fail to be symmetric in any useful sense, one would not get far by working with, say, $\mathbb{P} \circ (B, W^1, \widehat{\mu}^n(\tau^n), \tau_1^n)^{-1}$, which corresponds to *arbitrarily* choosing agent 1 as the representative. The same idea appears in the following converse to Theorem 4.4, which is an analog of Theorem 3.4 for the case of weak equilibria.

Theorem 4.5 Suppose assumption **D** holds. Let P be a weak MFE. Then there exist $\epsilon_n \rightarrow 0$ and ϵ_n -Nash equilibria $\tau^n = (\tau_1^n, \dots, \tau_n^n)$ such that

$$P = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{P} \circ \left(B, W^i, \widehat{\mu}^n(\tau^n), \tau_i^n \right)^{-1}.$$

In fact, if $\tau^* = \tau^*(B, W)$ is a strong MFE in the sense of Definition 4.1, then we can take τ^n of the form $\tau_i^n = \tau^*(B, W^i)$.

Finally, we state an existence result for weak MFE. Combined with Theorem 4.5, it shows that approximate n -player equilibria exist for the n -player games.

Theorem 4.6 Under assumption **D**, there exists a weak MFE.

Some comments are in order at this stage. Combining the two limit theorems tells us that the set of weak MFEs is precisely the set of limits of n -player approximate equilibria. If we can find a strong MFE τ^* , the converse limit theorem 4.5 shows how to construct from it an approximate n -player equilibria in a pleasantly symmetric and distributed form, as in Theorem 3.4. The general structure of the results and arguments are similar to [12, 28].

5 Proofs in the General Setup

This section proves the results of Sect. 3. Throughout this section, we work on the space $(\overline{\Omega}, \overline{\mathcal{F}}, \overline{\mathbb{F}}, \overline{\mathbb{P}})$ defined in Sect. 3.1. With some abuse of notation, any function ϕ on $\Omega^{\text{com}} \times \Omega^{\text{ind}}$ is automatically extended to all of $\overline{\Omega}$ by setting

$$\phi(\omega^0, \omega^1, \dots) := \phi(\omega^0, \omega^1).$$

Proof of Theorem 3.4

Abbreviate $\tau^{n,-k} := (\tau^n)^{-k}$, and define

$$\epsilon_n = \sup_{\tilde{\tau}} \mathbb{E} \left[F(W^0, W^1, \bar{\mu}^n(\tau^{n,-1}, \tilde{\tau}), \tilde{\tau}) \right] - \mathbb{E} \left[F(W^0, W^1, \bar{\mu}^n(\tau^n, \tau^1)) \right],$$

where the supremum is over $\mathbb{F}^1$ -stopping times. Clearly $\epsilon_n \geq 0$. By symmetry, the index 1 could be replaced by any $k \in \{1, \dots, n\}$. Hence, τ^n is an ϵ_n -Nash equilibrium for each n . We must only show that $\epsilon_n \rightarrow 0$.

First we show that

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\sup_{t \in \mathbb{T}} \left| F(W^0, W^1, \bar{\mu}^n(\tau^n, t) - F(W^0, W^1, \mu, t) \right| \right] = 0. \quad (5.1)$$

Note that the basic open sets of $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ are of the form

$$U = \left\{ m \in \mathcal{P}(\mathbb{T}) : \left| \int_{\mathbb{T}} \varphi_i d(m - \tilde{m}) \right| < \epsilon, i = 1, \dots, k \right\},$$

for $k \geq 1$, $\epsilon > 0$, and $\varphi_1, \dots, \varphi_k \in B(\mathbb{T})$. Because $(\tau^k = \tau^*(W^0, W^k))_{k=1}^\infty$ are conditionally i.i.d. given W^0 , and their common conditional law is $\mu = \mu(W^0)$, the law of large numbers yields

$$\mathbb{P} \left(\bar{\mu}^n(\tau^n) \notin U \mid W^0 = \omega^0 \right) \rightarrow 0, \quad (5.2)$$

for almost every ω^0 , for every basic $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ -open neighborhood U of $\mu(\omega^0)$. Thanks to the continuity assumption (A.2), for each $\delta > 0$ and almost every (ω^0, ω^1) we can find a basic $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ -open neighborhood U of $\mu(\omega^0)$ such that $v \in U$ implies

$$\sup_{t \in \mathbb{T}} \left| F(\omega^0, \omega^1, v, t) - F(\omega^0, \omega^1, \mu(\omega^0), t) \right| < \delta.$$

Thus, for a.e. ω^0 ,

$$\mathbb{P} \left(\sup_{t \in \mathbb{T}} \left| F(W^0, W^1, \bar{\mu}^n(\tau^n), t) - F(W^0, W^1, \mu(W^0), t) \right| \geq \delta \mid W^0 = \omega^0 \right) \rightarrow 0.$$

Thanks to assumption (A.3), the limit (5.1) follows from dominated convergence.

Next, we argue that

$$\lim_{n \rightarrow \infty} \sup_{\tilde{\tau}} \mathbb{E} \left[F(W^0, W^1, \bar{\mu}^n(\tau^{n,-1}, \tilde{\tau}), \tilde{\tau}) \right] = \sup_{\tilde{\tau}} \mathbb{E} \left[F(W^0, W^1, \mu, \tilde{\tau}) \right]. \quad (5.3)$$

Indeed, using the easy estimate

$$\sup_{t \in \mathbb{T}} \|\bar{\mu}^n(\tau^{n,-1}, t) - \bar{\mu}^n(\tau^n)\|_{TV} \leq 2/n,$$

along with (5.2), we deduce that for almost every ω^0 and for every basic $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ -open neighborhood U of $\mu(\omega^0)$, we have

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\bar{\mu}^n(\tau^{n,-1}, t) \notin U, \text{ for some } t \in \mathbb{T} \mid W^0 = \omega^0 \right) = 0.$$

Repeat the argument leading to (5.1) above to get

$$\mathbb{E} \left[\sup_{t \in \mathbb{T}} \left| F(W^0, W^1, \bar{\mu}^n(\tau^{n,-1}, t), t) - F(W^0, W^1, \mu, t) \right| \right] \rightarrow 0.$$

Finally, we conclude from (5.1), (5.3), and the optimality of τ^* that

$$\lim_{n \rightarrow \infty} \epsilon_n = \sup_{\tilde{\tau}} \mathbb{E} \left[F(W^0, W^1, \mu, \tilde{\tau}) \right] - \mathbb{E} \left[F(W^0, W^1, \mu, \tau^*) \right] = 0. \quad \square$$

Proof of Lemma 3.3

Fix $\epsilon > 0$. Find a finite set $0 = t_0 < t_1 < \dots < t_N = \sup \mathbb{T}$ such that $m(t_k, t_{k+1}) \leq \epsilon$ for every $k = 0, \dots, N-1$. Consider the $\sigma(\mathcal{P}(\mathbb{T}), B(\mathbb{T}))$ -open neighborhood U of m_0 given by

$$U = \{m \in \mathcal{P}(\mathbb{T}) : |m(t_k, t_{k+1}) - m_0(t_k, t_{k+1})| \vee |m[0, t_k] - m_0[0, t_k]| < \epsilon, \\ \forall k = 0, \dots, N-1\}.$$

For $m \in U$ and $t \in (t_k, t_{k+1})$ we have

$$\begin{aligned} |m[0, t] - m_0[0, t]| &\leq |m[0, t] - m[0, t_k]| + |m[0, t_k] - m_0[0, t_k]| + |m_0[0, t_k] - m_0[0, t]| \\ &\leq m(t_k, t) + m_0(t_k, t) + |m[0, t_k] - m_0[0, t_k]| \\ &\leq m(t_k, t_{k+1}) + m_0(t_k, t_{k+1}) + |m[0, t_k] - m_0[0, t_k]| \\ &\leq 3\epsilon. \end{aligned}$$

Setting $\pi = \{t_0, \dots, t_{N-1}\}$, we have, for $m \in U$,

$$\sup_{t \in \mathbb{T}} |m[0, t] - m_0[0, t]| = \sup_{t \notin \pi} |m[0, t] - m_0[0, t]| \vee \max_{t \in \pi} |m[0, t] - m_0[0, t]| \leq 3\epsilon.$$

□

Existence Under Supermodularity

In this section, we prove Theorem 3.5. Let $\mathcal{S}$ denote the set of (equivalence classes of a.s. equal) $\mathbb{F}^{\text{sig}}$ -stopping times, and let $\mathcal{M}$ denote the set of (equivalence classes of a.s. equal) $\mathcal{P}(\mathbb{T})$ -valued random variables μ , which are $\mathbb{F}^{\text{com}}$ -adapted in the sense that $\mu[0, t]$ is a.s. $\mathcal{F}_t^{\text{com}}$ -measurable for each t . Equip $\mathcal{S}$ with the almost sure partial order, meaning that we interpret the inequality $\tau \leq \tau'$ as holding almost surely. Equip $\mathcal{M}$ with the almost sure stochastic order, meaning that $\mu' \geq \mu$ if and only if $\mu'[0, t] \leq \mu[0, t]$ a.s. for each $t \in [0, T]$, and note that right-continuity renders the order of quantifiers inconsequential. Note that $\mathcal{M}$ is a lattice, namely a partially ordered set in which every two elements have a unique least upper bound and a unique greatest lower bound; for example $\mu \vee \mu'$ is the random measure defined by $(\mu \vee \mu')[0, t] = \mu[0, t] \wedge \mu'[0, t]$. On the other hand, $\mathcal{S}$ is a *complete* lattice in the sense that it is a partially ordered set in which every subset has both a supremum and an infimum. Indeed, the notion of “essential supremum” provides the correct supremum operation on $\mathcal{S}$, and the completeness follows from the assumption (B.1) that the filtration $\mathbb{F}^{\text{sig}}$ is right-continuous. The completeness of the lattice of stopping times is surely known, but we prove it in the Appendix (Theorem B.2) as we were unable to locate a precise reference.

Now define $J : \mathcal{M} \times \mathcal{S} \rightarrow \mathbb{R}$ by

$$J(\mu, \tau) = \mathbb{E}[F(B, W, \mu, \tau)].$$

Note that $J(\mu, \tau)$ is trivially supermodular in τ , in the sense that

$$J(\mu, \tau \vee \tau') + J(\mu, \tau \wedge \tau') \geq J(\mu, \tau) + J(\mu, \tau'),$$

for every $\mu \in \mathcal{M}$ and every pair $\tau, \tau' \in \mathcal{S}$. In fact, this holds with equality, which follows from taking expectations on both sides of the identity

$$F(B, W, \mu, \tau \vee \tau') + F(B, W, \mu, \tau \wedge \tau') = F(B, W, \mu, \tau) + F(B, W, \mu, \tau').$$

Moreover, assumption (B.2) ensures that J has increasing differences with respect to μ , in the sense that

$$J(\mu', \tau') - J(\mu', \tau) \geq J(\mu, \tau') - J(\mu, \tau),$$

whenever $\tau, \tau' \in \mathcal{S}$ and $\mu, \mu' \in \mathcal{M}$ satisfy $\tau \leq \tau'$ and $\mu \leq \mu'$. From Topkis’s monotonicity theorem [30], we deduce that the set-valued map

$$\Phi(\mu) := \arg \max_{\tau \in \mathcal{S}} J(\mu, \tau)$$

is increasing in the strong set order, meaning that whenever $\mu, \mu' \in \mathcal{M}$ satisfy $\mu \leq \mu'$, and whenever $\tau \in \Phi(\mu)$ and $\tau' \in \Phi(\mu')$, we have $\tau \vee \tau' \in \Phi(\mu')$ and $\tau \wedge \tau' \in \Phi(\mu)$. It is readily checked that J is order upper semicontinuous in τ , using the assumed

upper semicontinuity of F in τ along with Fatou's lemma, justified by the integrability assumption (3.3). By [30, Theorem 1], this implies that for every μ , $\Phi(\mu)$ is a nonempty complete lattice. In particular, it has a maximum, which we denote $\phi^*(\mu)$ and a minimum which we denote by $\phi_*(\mu)$. Note that $\phi^* : \mathcal{M} \rightarrow \mathcal{S}$ is increasing in the sense that $\mu' \geq \mu$ implies $\phi^*(\mu') \geq \phi^*(\mu)$. Moreover, it is plain to check that the function $\psi : \mathcal{S} \rightarrow \mathcal{M}$ defined by $\psi(\tau) = \text{Law}(\tau|W^0)$ is monotone. Thus $\phi^* \circ \psi$ is a monotone map from $\mathcal{S}$ to itself, and since $\mathcal{S}$ is a complete lattice we conclude from Tarski's fixed point theorem that there exists τ such that $\tau = \phi^*(\psi(\tau))$. It is readily verified that any such fixed point τ is a strong MFE, in the sense of Definition 4.1.

Now, under the additional assumption C, we prove the last claim in the statement of Theorem 3.5. Define $\tau_0 \equiv \sup \mathbb{T}$, and by induction $\tau_n = \phi^* \circ \psi(\tau_{n-1})$ for $n \geq 1$. Clearly, $\tau_1 \leq \tau_0$. Now if we assume $\tau_n \leq \tau_{n-1}$, then the monotonicity of $\phi^* \circ \psi$ proved earlier implies $\tau_{n+1} = \phi^* \circ \psi(\tau_n) \leq \phi^* \circ \psi(\tau_{n-1}) = \tau_n$. If we define τ^* as the a.s. limit of the nonincreasing sequence $(\tau_n)_{n \geq 1}$ of stopping times, then $\tau^* \in \mathcal{S}$ because the lattice $\mathcal{S}$ is complete (see Theorem B.2). We claim that τ^* is a MFE. Note that $\psi(\tau_n) \rightarrow \psi(\tau^*)$ weakly almost surely, because $\tau_n \rightarrow \tau^*$. The assumption (C.1) of joint continuity of $F = F(\omega^0, \omega^1, m, t)$ in (m, t) , along with the uniform integrability assumption (C.2) ensure by dominated convergence that

$$J(\psi(\tau^*), \tau^*) = \lim_{n \rightarrow \infty} J(\psi(\tau_n), \tau_{n+1}).$$

Moreover, for any $\sigma \in \mathcal{S}$, the fact that $\tau_{n+1} \in \Phi(\psi(\tau_n))$ implies

$$J(\psi(\tau_n), \tau_{n+1}) \geq J(\psi(\tau_n), \sigma).$$

Pass to the limit on both sides to get

$$J(\psi(\tau^*), \tau^*) \geq J(\psi(\tau^*), \sigma).$$

This shows $\tau^* \in \Phi(\psi(\tau^*))$, and in particular τ^* is a MFE.

Similarly, define $\theta_0 \equiv 0$, and by induction $\theta_n = \phi_* \circ \psi(\theta_{n-1})$ for $n \geq 1$. Clearly, $\theta_1 \geq \theta_0$, and as above, we prove by induction that $\theta_n \geq \theta_{n-1}$. Next, we define θ^* as the a.s. limit of the nondecreasing sequence $(\theta_n)_{n \geq 1}$ of stopping times. Conclude as before that θ^* is a MFE.

Finally, it is plain to check that if τ is any MFE, it is a fixed point of the set-valued map $\Phi \circ \psi$, in the sense that $\tau \in \Phi(\psi(\tau))$. Trivially, $\theta_0 = 0 \leq \tau \leq \sup \mathbb{T} = \tau_0$. Applying $\phi_* \circ \psi$ and $\phi^* \circ \psi$ repeatedly to the left and right sides, respectively, we conclude that $\theta_n \leq \tau \leq \tau_n$ for each n , and thus $\theta^* \leq \tau \leq \tau^*$. $\square$

Remark 5.1 The above proof shows that, under the full continuity assumption, there is no need to use Tarski's theorem to prove existence, as the MFEs τ^* and θ^* are constructed inductively.

6 Compatibility and the Density of Non-randomized Stopping Times

This section elaborates on the crucial notion of *compatibility* introduced in property (3) of Definition 4.2 and, in doing so, takes some first steps toward proving the results announced in Sect. 4. Here, our goal is to discuss some important facts about these compatibility properties, namely how to approximate compatible (randomized) stopping times with nonrandomized stopping times. Essentially, this has to do with filtration enlargements. To say that $\mathcal{F}_{t+}^{\tau}$ is conditionally independent of $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_{t+}^{B,W,\mu}$ is the same as saying that $\mathcal{F}_{t+}^{B,W,\mu,\tau}$ is conditionally independent of $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_{t+}^{B,W,\mu}$. To say that this holds for every $t \in [0, T)$, it turns out, is equivalent to saying that every $\mathbb{F}_+^{B,W,\mu}$ -martingale remains a $\mathbb{F}_+^{B,W,\mu,\tau}$ -martingale. Many different names and characterizations are associated to this property of a filtration enlargement, such as the H-hypothesis [8], immersion [24], very good extensions [23], and natural extensions [15], while we borrow the term *compatible* from Kurtz [25], to be consistent with other works on mean field games [11, 12, 28]. Before we proceed, we recall a useful result on weak convergence which will be used repeatedly:

Lemma 6.1 (Corollary 2.9 and Theorem 2.16 of [22]) *Suppose E and E' are Polish spaces. Suppose $P_n, P \in \mathcal{P}(E \times E')$ satisfy $P_n \rightarrow P$, and suppose that every P_n has the same E -marginal. That is, $P_n(\cdot \times E')$ does not depend on n . Then, for every bounded measurable function $\phi : E \times E' \rightarrow \mathbb{R}$ such that $\phi(x, \cdot)$ is continuous on E' for μ -almost every $x \in E$, we have*

$$\int \phi dP_n \rightarrow \int \phi dP.$$

Of utmost importance to us is the behavior of compatibility under weak limits of the underlying probability measures. The key result is the following, which says roughly that a compatible process is the weak limit of adapted processes.

Proposition 6.2 *Let $\mathcal{X}$ and $\mathcal{Y}$ be Polish spaces, with $\mathcal{X}$ homeomorphic to a convex subset of a locally convex space. Let $Y = (Y_1, \dots, Y_N)$ and $X = (X_1, \dots, X_N)$ be $\mathcal{Y}$ - and $\mathcal{X}$ -valued stochastic processes, respectively, defined on a common probability space. For $R \in \{X, Y\}$, let $\mathcal{F}_n^R = \sigma\{R_1, \dots, R_n\}$ denote the filtration generated by R . Assume that the law of Y_1 is nonatomic. Lastly, assume that X is compatible with Y in the sense that $\mathcal{F}_n^X$ is conditionally independent of $\mathcal{F}_N^Y$ given $\mathcal{F}_n^Y$, for each $n = 1, \dots, N$. Then there exist continuous functions $h_k^j : \mathcal{Y}^k \rightarrow \mathcal{X}$, for $k \in \{1, \dots, N\}$ and $j \geq 1$, such that*

$$(Y, (h_1^j(Y_1), h_2^j(Y_1, Y_2), \dots, h_N^j(Y_1, \dots, Y_N))) \rightarrow (Y, X)$$

in law in the space $\mathcal{Y}^N \times \mathcal{X}^N$, as $j \rightarrow \infty$. In particular, there exist Y -adapted $\mathcal{X}$ -valued processes $X^j = (X_1^j, \dots, X_N^j)$ such that $(Y, X^j) \Rightarrow (Y, X)$.

Proof See Appendix A. □

6.1 Randomized Stopping Times

This section is devoted to some compactness properties of randomized stopping times, analogous to, but extending results of Baxter and Chacon [7]. Abbreviate

$$\Omega_{\text{input}} = \mathcal{C}^2 \times \mathcal{P}(\mathcal{C} \times [0, T]).$$

For this section, fix a measure $\rho \in \mathcal{P}(\Omega_{\text{input}})$, to represent a joint law of (B, W, μ) , and assume, under ρ , that (B, W) are Wiener processes with respect to the filtration $\mathbb{F}^{B, W, \mu}$ (and thus also with respect to the right-filtration $\mathbb{F}_+^{B, W, \mu}$). Note that $\Omega = \Omega_{\text{input}} \times [0, T]$.

We next define three sets of probability measures on Ω , corresponding to various notions of (randomized) stopping times:

- $\mathcal{R}^+(\rho)$ is the set of joint laws $P \in \mathcal{P}(\Omega)$ with Ω_{input} -marginal ρ such that $\mathcal{F}_{t+}^\tau$ is conditionally independent of (B, W, μ) given $\mathcal{F}_{t+}^{B, W, \mu}$ for every $t \in [0, T)$. That is, $\mathcal{R}^+(\rho)$ is the set of P satisfying the compatibility property (3) of Definition 4.2.
- $\mathcal{R}(\rho)$ is the set of joint laws $P \in \mathcal{P}(\Omega)$ with Ω_{input} -marginal ρ such that $\mathcal{F}_t^\tau$ is conditionally independent of (B, W, μ) given $\mathcal{F}_t^{B, W, \mu}$, for every $t \in [0, T)$.
- $\mathcal{R}_0(\rho)$ is the set of $P \in \mathcal{P}(\Omega)$ with Ω_{input} -marginal ρ under which τ is a stopping time relative to the P -completion of $\mathbb{F}^{B, W, \mu}$, and moreover τ is of the form $\tau = \hat{\tau}(B, W, \mu)$ for some continuous function $\hat{\tau} : \Omega_{\text{input}} \rightarrow [0, T]$.

Both sets $\mathcal{R}^+(\rho)$ and $\mathcal{R}(\rho)$ represent slightly different notions of randomized stopping time, though we will soon see that $\mathcal{R}^+(\rho) = \mathcal{R}(\rho)$. On the other hand, $\mathcal{R}_0(\rho)$ should be seen as the set of (joint laws of) bona fide $\mathbb{F}^{B, W, \mu}$ -stopping times, with the useful additional property that τ can be written as a continuous function of (B, W, μ) .

Remark 6.3 Suppose $\rho(db, dw, dm) = \mathcal{W}(db)\mathcal{W}(dw)\delta_{\hat{m}(b)}(dm)$ for some measurable function $\hat{m} : \mathcal{C} \rightarrow \mathcal{P}(\mathcal{C} \times [0, T])$. Suppose $\hat{m}$ is adapted in the sense that $b \mapsto \hat{m}(b)(C)$ is $\mathcal{F}_t^B$ -measurable whenever $C \in \mathcal{F}_t^{W, \tau}$, for $t \in [0, T]$. Then, under ρ , $\mathcal{F}_t^{B, W, \mu} = \mathcal{F}_t^{B, W}$ a.s., for every t . It is then easy to argue that $\mathcal{R}^+(\rho)$ (resp. $\mathcal{R}(\rho)$) is precisely the set of joint laws $P \in \mathcal{P}(\Omega)$ with Ω_{input} -marginal ρ such that (B, W) is a Wiener process with respect to the full filtration $\mathbb{F}_+^{B, W, \mu, \tau}$ (resp. $\mathbb{F}^{B, W, \mu, \tau}$). Indeed, because $\mathcal{F}_T^{B, W}$ can be split into two independent parts, $\mathcal{F}_T^{B, W} = \sigma\{(B_s - B_t, W_s - W_t) : s \in [t, T]\} \vee \mathcal{F}_t^{B, W}$, it holds that $\mathcal{F}_t^\tau$ is conditionally independent of $\mathcal{F}_t^{B, W}$ given $\mathcal{F}_t^{B, W}$ if and only if $\mathcal{F}_t^{B, W, \tau}$ is independent of $\sigma\{(B_s - B_t, W_s - W_t) : s \in [t, T]\}$.

Theorem 6.4 For ρ as above, $\mathcal{R}^+(\rho)$ is convex and compact and equals the closure of $\mathcal{R}_0(\rho)$. Moreover, $\mathcal{R}^+(\rho) = \mathcal{R}(\rho)$.

Before turning to the proof, we state an immensely useful corollary:

Corollary 6.5 Assume that F is bounded and jointly measurable and that $t \mapsto F(b, w, m, t)$ is continuous, for every m and $\mathcal{W}^2$ -almost every (b, w) . For ρ as above, we have

$$\sup_{P \in \mathcal{R}^+(\rho)} \mathbb{E}^P[F(B, W, \mu^\tau, \tau)] = \sup_{P \in \mathcal{R}_0(\rho)} \mathbb{E}^P[F(B, W, \mu^\tau, \tau)].$$

Proof By Theorem 6.4, $\mathcal{R}_0(\rho)$ is dense in $\mathcal{R}^+(\rho)$. It suffices to show that $P \mapsto \mathbb{E}^P[F(B, W, \mu, \tau)]$ is continuous on $\mathcal{R}^+(\rho)$. But this follows from the assumption on F and from Lemma 6.1. $\square$

We precede the proof of Theorem 6.4 with a preparatory lemma, which allows us to map continuously between stopping times and càdlàg processes of a certain form. In the rest of this section, let $D = D([0, T]; \mathbb{R}_+)$ denote the set of càdlàg functions (i.e., right-continuous with left limits) functions from $[0, T]$ to $\mathbb{R}_+ = [0, \infty)$. Endow D with the usual Skorohod J_1 topology. As usual, for $h \in D$, write $h(t-) = \lim_{s \uparrow t} h(s)$ for $t \in (0, T]$ and $h(0-) = h(0)$.

Lemma 6.6 Define $\Phi : D \rightarrow [0, T]$ by

$$\Phi(h) = \inf \{t \geq 0 : h(t) \geq 1/2\} \wedge T.$$

Let $\mathcal{S}$ denote the set of nondecreasing $h \in D$ for which $h(t-) \leq 1/2 \leq h(t)$ implies $t = \Phi(h)$. Then Φ is continuous at each point of $\mathcal{S}$.

Proof Let $h_n \rightarrow h$ in D , where $h \in \mathcal{S}$. Let $s_n = \Phi(h_n)$, and note that $(s_n)_{n=1}^\infty$ is bounded. Let $(s_{n_k})_{k=1}^\infty$ denote any convergent subsequence, and let $s \in [0, T]$ denote its limit. Assume first that $0 < s < T$, so that without loss of generality we may take $0 < s_{n_k} < T$ for every k . Because $h_{n_k} \rightarrow h$ and $s_{n_k} \rightarrow s$, it follows that $(h_{n_k}(s_{n_k}))_{k=1}^\infty$ is bounded, and its limit points are contained in $\{h(s-), h(s)\}$ (see [16, Proposition 3.6.5]). Because $h_{n_k}(s_{n_k}) \geq 1/2$ for every k and $h(s) \geq h(s-)$ (as $h \in \mathcal{S}$), we conclude that $h(s) \geq 1/2$. On the other hand, for $\epsilon > 0$, $(h_{n_k}(s_{n_k} - \epsilon))_{k=1}^\infty$ is bounded, and its limit points are contained in $\{h((s - \epsilon)-), h(s - \epsilon)\}$. Because $h_{n_k}(s_{n_k} - \epsilon) < 1/2$ for every k and $h((s - \epsilon)-) \leq h(s - \epsilon)$, we conclude that $h(s - \epsilon) \leq 1/2$. Sending $\epsilon \downarrow 0$ yields $h(s-) \leq 1/2 \leq h(s)$ and thus $s = \Phi(h)$.

Next, suppose $s = T$. Then again $(h_{n_k}(s_{n_k} - \epsilon))_{k=1}^\infty$ is bounded, and its limit points are contained in $\{h((T - \epsilon)-), h(T - \epsilon)\}$. Because $h_{n_k}(s_{n_k} - \epsilon) < 1/2$ for each k and $h((T - \epsilon)-) \leq h(T - \epsilon)$, we conclude that $h((T - \epsilon)-) \leq 1/2$. This implies $h(T-) \leq 1/2$, which is enough to show that $\Phi(h) = T$; indeed, either $h(T) \geq 1/2$, in which case $\Phi(h) = T$ because $h \in \mathcal{S}$, or $h(T) < 1/2$, in which case $h(t) < 1/2$ for all $t \in [0, T]$ and again $\Phi(h) = T$.

Finally, suppose $s = 0$. Then $h_{n_k}(s_{n_k}) \rightarrow h(0) = h(0-)$, which shows $h(0) \geq 1/2$. Thus $\Phi(h) = 0$. $\square$

Before we begin the proof of Theorem 6.4, notice that $\mathcal{F}_t^{B, W, \mu} = \sigma\{B_{\cdot \wedge t}, W_{\cdot \wedge t}, \mu^t\}$, where m^t denotes the image of a measure $m \in \mathcal{P}(\mathcal{C} \times [0, T])$ through the map $\mathcal{C} \times [0, T] \ni (w, s) \mapsto (w_{\cdot \wedge t}, s \wedge t)$. This makes it clear that $\mathcal{F}_t^{B, W, \mu}$ is generated by the continuous $\mathcal{F}_t^{B, W, \mu}$ -measurable functions. Similarly, $\mathcal{F}_t^\tau = \sigma\{\tau \wedge t\}$ is generated by the continuous $\mathcal{F}_t^\tau$ -measurable functions.

Proof of Theorem 6.4

We break the proof up into four claims:

$\mathcal{R}(\rho)$ is compact Because the Ω_{input} -marginal of any element of $\mathcal{R}(\rho)$ is ρ , and because $[0, T]$ is compact, it is immediate that $\mathcal{R}(\rho)$ is tight. To show $\mathcal{R}(\rho)$ is closed, let $P_n \rightarrow P$ in $\mathcal{P}(\Omega_{\text{input}})$, with $P_n \in \mathcal{R}(\rho)$. Let f_t , g_T , and g_t be bounded continuous functions on $[0, T]$, Ω_{input} , and Ω_{input} , respectively, and assume they are measurable with respect to $\mathcal{F}_t^\tau$, $\mathcal{F}_t^{B,W,\mu}$, and $\mathcal{F}_t^{B,W,\mu}$. Find a bounded $\mathcal{F}_t^{B,W,\mu}$ -measurable function ϕ_t on Ω_{input} such that $\phi_t(B, W, \mu) = \mathbb{E}^P[g_T(B, W, \mu)|\mathcal{F}_t^{B,W,\mu}]$. Because $P_n \circ (B, W, \mu)^{-1} = P \circ (B, W, \mu)^{-1} = \rho$ for every n , we have $\phi_t(B, W, \mu) = \mathbb{E}^{P_n}[g_T(B, W, \mu)|\mathcal{F}_t^{B,W,\mu}]$. Thus, by Lemma 6.1,

$$\begin{aligned} \mathbb{E}^P[f_t(\tau)g_T(B, W, \mu)g_t(B, W, \mu)] &= \lim_{n \rightarrow \infty} \mathbb{E}^{P_n}[f_t(\tau)g_T(B, W, \mu)g_t(B, W, \mu)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{P_n}[f_t(\tau)\phi_t(B, W, \mu)g_t(B, W, \mu)] \\ &= \mathbb{E}^P[f_t(\tau)\phi_t(B, W, \mu)g_t(B, W, \mu)]. \end{aligned}$$

As remarked before the proof, the continuous bounded functions generate $\mathcal{F}_t^{B,W,\mu}$ and $\mathcal{F}_t^\tau$, and we conclude that

$$\begin{aligned} \mathbb{E}^P[f_t(\tau)g_T(B, W, \mu)g_t(B, W, \mu)] \\ = \mathbb{E}^P[f_t(\tau)\mathbb{E}^P[g_T(B, W, \mu)|\mathcal{F}_t^{B,W,\mu}]g_t(B, W, \mu)], \end{aligned}$$

for all bounded functions f_t , g_T , g_t with the same measurability requirements as above, but without the continuity requirements. This shows that $\mathcal{F}_t^\tau$ is conditionally independent of $\mathcal{F}_t^{B,W,\mu}$ given $\mathcal{F}_t^{B,W,\mu}$ under P , for every $t \in [0, T]$, so $P \in \mathcal{R}(\rho)$.

$\mathcal{R}(\rho) = \mathcal{R}^+(\rho)$ First we show $\mathcal{R}(\rho) \subset \mathcal{R}^+(\rho)$. Fix $t \in [0, T]$ and $P \in \mathcal{R}(\rho)$. For $A \in \mathcal{F}_t^{B,W,\mu}$ and $C \in \mathcal{F}_t^\tau$, we have

$$\mathbb{P}(C|\mathcal{F}_t^{B,W,\mu})\mathbb{P}(A|\mathcal{F}_t^{B,W,\mu}) = \mathbb{P}(C \cap A|\mathcal{F}_t^{B,W,\mu}).$$

By backward martingale convergence, taking decreasing limits in t yields

$$\mathbb{P}(C|\mathcal{F}_{t+}^{B,W,\mu})\mathbb{P}(A|\mathcal{F}_{t+}^{B,W,\mu}) = \mathbb{P}(C \cap A|\mathcal{F}_{t+}^{B,W,\mu}).$$

This shows that $\mathcal{R}(\rho) \subset \mathcal{R}^+(\rho)$, and we know from before that $\mathcal{R}(\rho)$ is closed. Hence, it suffices to show that every point $P \in \mathcal{R}^+(\rho)$ is the limit point of a sequence in $\mathcal{R}(\rho)$. To see this, set

$$P_n := P \circ (B, W, \mu, (\tau + 1/n) \wedge T)^{-1},$$

which converges weakly to P . Because $\mathcal{F}_{t+}^{\tau}$ is conditionally independent of $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_{t+}^{B,W,\mu}$ under P , we have, for $0 < s \leq t < T$,

$$\begin{aligned} P_n(\tau \leq s | \mathcal{F}_T^{B,W,\mu}) &= P(\tau \leq s - 1/n | \mathcal{F}_T^{B,W,\mu}) = P(\tau \leq s - 1/n | \mathcal{F}_{(t-1/n)+}^{B,W,\mu}) \\ &= P_n(\tau \leq s | \mathcal{F}_{(t-1/n)+}^{B,W,\mu}). \end{aligned}$$

Because $\mathcal{F}_{(t-1/n)+}^{B,W,\mu} \subset \mathcal{F}_t^{B,W,\mu}$, we conclude that

$$P_n(\tau \leq s | \mathcal{F}_T^{B,W,\mu}) = P_n(\tau \leq s | \mathcal{F}_t^{B,W,\mu}). \quad (6.1)$$

Since $P_n(\tau = 0) = 0$, we also have (6.1) for $s = 0$. Since (6.1) holds for all $s \in [0, t]$, we conclude that $\mathcal{F}_{t+}^{\tau} = \sigma\{\{\tau \leq s\} : s \leq t\}$ is independent of $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_t^{B,W,\mu}$ under P_n , for $t \in (0, T)$. To conclude that $P_n \in \mathcal{R}(\rho)$, we must still check that $\mathcal{F}_0^{\tau}$ is conditionally independent of $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_0^{B,W,\mu}$ under P_n . But this is obvious, as $\mathcal{F}_0^{\tau} = \sigma\{\tau \wedge 0\}$ is the trivial σ -field.

$\mathcal{R}(\rho)$ is convex To check that $\mathcal{R}(\rho)$ is convex, note that $\mathcal{R}(\rho)$ is the set of $P \in \mathcal{P}(\Omega_{\text{input}} \times [0, T])$ with first marginal ρ for which

$$\begin{aligned} &\mathbb{E}^P [\phi_t(\tau) \psi(B, W, \mu) \psi_t(B, W, \mu)] \\ &= \mathbb{E}^P \left[\mathbb{E}^P \left[\phi_t(\tau) | \mathcal{F}_t^{B,W,\mu} \right] \psi(B, W, \mu) \psi_t(B, W, \mu) \right], \end{aligned}$$

for every $t \in [0, T]$ and every triple of bounded functions ϕ_t , ψ , and ψ_t , measurable with respect to $\mathcal{F}_t^{\tau}$, $\mathcal{F}_T^{B,W,\mu}$, and $\mathcal{F}_t^{B,W,\mu}$, respectively. Disintegrate any $P \in \mathcal{R}(\rho)$ by writing $P(d\omega, ds) = \rho(d\omega)P(\omega, ds)$, and note that the above equation is equivalent to

$$\begin{aligned} &\int_{\Omega_{\text{input}} \times [0, T]} P(d\omega, du) \psi(\omega) \psi_t(\omega) \phi_t(u) \\ &= \int_{\Omega_{\text{input}}} \rho(d\omega) \psi(\omega) \psi_t(\omega) \int_{[0, T]} P(\omega, du) \phi_t(u). \end{aligned}$$

This is clearly a convex constraint on P , and we conclude that $\mathcal{R}(\rho)$ is convex.

$\mathcal{R}^+(\rho)$ is contained in the closure of $\mathcal{R}_0(\rho)$ Let $P \in \mathcal{R}^+(\rho)$, and consider the process $H_t = 1_{\{\tau \leq t\}}$ defined on $\Omega_{\text{input}} \times [0, T]$. Note that H is $\mathbb{F}_+^{\tau}$ -adapted, because $\{H_t = 1\} = \{\tau \leq t\} \in \mathcal{F}_{t+}^{\tau}$. As a first approximation, note that the càdlàg process $H_t^n = 1_{\{(\tau+1/n) \wedge T \leq t\}}$ is $\mathbb{F}^{\tau}$ -adapted and converges almost surely to H in the Skorohod topology. Because $\Phi(H^n) = (\tau + 1/n) \wedge T \rightarrow \tau = \Phi(H)$ a.s. by Lemma 6.6, we may henceforth assume that H is in fact $\mathbb{F}^{\tau}$ -adapted.

Next, by a routine approximation we may find a sequence of càdlàg $\mathbb{F}^{\tau}$ -adapted processes H^n converging almost surely to H and of the form

$$H_t^n = \sum_{k=1}^K h_k^n 1_{[t_k^n, t_{k+1}^n)}(t),$$

where $0 < t_1^n < t_2^n < \dots < t_K^n = T < t_{K+1}^n$, and where $h_k^n \geq 0$ is an $\mathcal{F}_{t_k^n}^\tau$ -measurable random variables. Replacing h_k^n by $\max_{j=1,\dots,k} h_j^n$ does not change the value of $\Phi(H^n)$, which again converges almost surely to $\Phi(\hat{H}) = \tau$ by Lemma 6.6, which applies because $H \in \mathcal{S}$ a.s. As a final approximation, let $\hat{H}_t^n = H_t^n + t/n$; this last approximation accounts for the fact that H^n may not belong a.s. to $\mathcal{S}$, while $\hat{H}$ does. Note that $|H_t^n - \hat{H}_t^n| \rightarrow 0$ uniformly in t , so $\lim_n \Phi(\hat{H}^n) = \lim_n \Phi(H^n) = \tau$. With these approximations, we may then assume henceforth that H itself is increasing and of the form

$$H_t = \sum_{k=1}^K h_k 1_{[t_k, t_{k+1})}(t) + \frac{t}{n},$$

where $0 < t_1 < t_2 < \dots < t_K = T < t_{K+1}$, and where h_k is $\mathcal{F}_{t_k}^\tau$ -measurable for each k .

Now define S_t for $t \in [0, T]$ by

$$S_t = (B_{\cdot \wedge t}, W_{\cdot \wedge t}, \mu^t),$$

where m^t was defined just before the start of the proof. Then $S = (S_t)_{t \in [0, T]}$ is a continuous $\mathbb{B}^{B, W, \mu}$ -adapted process, with values in Ω_{input} . In fact, $\mathcal{F}_t^{B, W, \mu} = \sigma(S_t) = \sigma(S_{\cdot \wedge t})$ for every t . For every $k = 1, \dots, K$, note that $(h_1, \dots, h_k)$ is conditionally independent of S given $S_{\cdot \wedge t_k}$. It follows from Proposition 6.2 that there exists a sequence of continuous functions $g_k^n : \Omega_{\text{input}} \rightarrow \mathbb{R}_+$ such that $g_k^n(S)$ is $\sigma(S_{t_k}) = \mathcal{F}_{t_k}^{B, W, \mu}$ -measurable for each k and

$$(S, g_1^n(S), \dots, g_K^n(S)) \Rightarrow (S, h_1, \dots, h_K), \text{ in } \Omega_{\text{input}} \times \mathbb{R}_+^K, \text{ as } n \rightarrow \infty.$$

Now define

$$H_t^n = H_t^n(S) = \sum_{k=1}^K g_k^n(S) 1_{[t_k, t_{k+1})}(t) + \frac{t}{n}.$$

It follows that $(S, H^n) \Rightarrow (S, H)$, and because H belongs almost surely to the set $\mathcal{S}$ of Lemma 6.6 we have $(S, \Phi(H^n)) \Rightarrow (S, \Phi(H)) = (S, \tau)$. Now let $\tilde{g}_k^n(s) = \max_{j=1,\dots,k} g_j^n(s)$, and define

$$\tilde{H}_t^n = \tilde{H}_t^n(S) = \sum_{k=1}^K \tilde{g}_k^n(S) 1_{[t_k, t_{k+1})}(t).$$

Then $\Phi(\tilde{H}^n) = \Phi(H^n)$ almost surely, so again we have $(S, \Phi(\tilde{H}^n)) \Rightarrow (S, \tau)$. Now, because each $\tilde{g}_k^n$ is continuous, we may view $\tilde{H}(\cdot)$ as a continuous map from Ω_{input} to the subset $\mathcal{S}$ of D defined in Lemma 6.6. Hence, Lemma 6.6 ensures that $\Phi(\tilde{H}^n(\cdot))$ is a continuous map from Ω_{input} to $[0, T]$. The law of $(S, \Phi(\tilde{H}^n))$ thus belongs to $\mathcal{R}_0(\rho)$. $\square$

Proof of Proposition 4.3

With Corollary 6.5 in hand, we are now ready to prove Proposition 4.3. It is readily checked that P satisfies property (2) of Definition 4.2 of a weak MFE. As τ^* is a $\mathbb{F}^{B,W}$ -stopping time, $\mathcal{F}_{t+}^\tau$ is contained in the P -completion of $\mathcal{F}_t^{B,W,\mu}$, and the compatibility property (3) holds easily (noting that compatibility is not sensitive to the completion of filtrations). To prove (1) is slightly more involved: Clearly $P \circ (B, W)^{-1} = \mathcal{W}^2$, where as usual $\mathcal{W}$ denotes Wiener measure. Note also that if $g_t : \mathcal{C} \rightarrow \mathbb{R}$ is bounded and $\mathcal{F}_t^{B,W}$ -measurable, then

$$\mathbb{E}[g_t(B, W)|B] = \int_{\mathcal{C}} g_t(B, w) \mathcal{W}(dw) = \mathbb{E}[g_t(B, W)|\mathcal{F}_t^B], \text{ a.s.}$$

Thus, since τ is a.s. (B, W) -measurable, if $t \in [0, T]$ and $C \in \mathcal{F}_t^{W,\tau}$ then

$$\mu(C) = P((W, \tau) \in C|B) = P((W, \tau) \in C|\mathcal{F}_t^B), \text{ } P - \text{a.s.}$$

This shows that (the completion of) $\mathcal{F}_t^\mu$ is contained in (the completion of) $\mathcal{F}_t^B$. Hence, under P , the completion of $\mathcal{F}_t^{B,W,\mu,\tau}$ is contained in that of $\mathcal{F}_t^{B,W}$, which proves property (1). The weak fixed point condition (5) holds because μ is B -measurable. Finally, the optimality condition (4) follows from Corollary 6.5. $\square$

6.2 A Shortcut to Compatibility

As a final preparatory step, before proving the main results we state one last useful lemma. It allows us to check a much simpler criterion in place of the compatibility property (3) of Definition 4.2, which does not behave too well under limits. In fact, this lemma is precisely the reason we work with the conditional *joint* law of (W, τ) and not just τ itself.

Lemma 6.7 *Suppose $P \in \mathcal{P}(\Omega)$ satisfies properties (2) and (5) of Definition 4.2. Suppose also that $\mathcal{F}_T^{B,\mu} \vee \mathcal{F}_t^{W,\tau}$ is independent of $\sigma\{W_s - W_t : s \in [t, T]\}$, for every $t \in [0, T)$. Then P satisfies property (3) of Definition 4.2; that is, $\mathcal{F}_{t+}^\tau$ is conditionally independent of $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_{t+}^{B,W,\mu}$, for every $t \in [0, T]$.*

Proof According to the final claim of Theorem 6.4, it suffices to check that $\mathcal{F}_t^\tau$ is conditionally independent of $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_t^{B,W,\mu}$, for every $t \in [0, T]$. Fix bounded functions f_t, g_T, g_t, h^+ , and h_t , such that $f_t(\tau)$ is $\mathcal{F}_t^\tau$ -measurable, $g_T(B, \mu)$ is $\mathcal{F}_T^{B,\mu}$ -measurable, $g_t(B, \mu)$ is $\mathcal{F}_t^{B,\mu}$ -measurable, $h^+(W)$ is $\sigma\{W_s - W_t : s \in [t, T]\}$ -measurable, and $h_t(W)$ is $\mathcal{F}_t^W$ -measurable. Compute

$$\begin{aligned} & \mathbb{E}[f_t(\tau)g_T(B, \mu)g_t(B, \mu)h^+(W)h_t(W)] \\ &= \mathbb{E}[f_t(\tau)g_T(B, \mu)g_t(B, \mu)h_t(W)]\mathbb{E}[h^+(W)] \\ &= \mathbb{E}\left[g_t(B, \mu)g_T(B, \mu) \int f_t(s)h_t(w)\mu(dw, ds)\right]\mathbb{E}[h^+(W)] \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E} \left[g_t(B, \mu) \mathbb{E} \left[g_T(B, \mu) | \mathcal{F}_t^{B, \mu} \right] \int f_t(s) h_t(w) \mu(dw, ds) \right] \mathbb{E}[h^+(W)] \\
 &= \mathbb{E} \left[g_t(B, \mu) h_t(W) f_t(\tau) \mathbb{E} \left[g_T(B, \mu) | \mathcal{F}_t^{B, \mu} \right] \right] \mathbb{E}[h^+(W)] \\
 &= \mathbb{E} \left[g_t(B, \mu) h_t(W) f_t(\tau) \mathbb{E} \left[g_T(B, \mu) h^+(W) | \mathcal{F}_t^{B, W, \mu} \right] \right].
 \end{aligned}$$

The first step used the assumed independence, whereas the second and fourth used the fixed point property $\mu = P((W, \tau) \in \cdot | B, \mu)$. The third step used the fact that $\int \phi d\mu$ is $\mathcal{F}_t^\mu$ -measurable (and thus $\mathcal{F}_t^{B, \mu}$ -measurable) for every bounded $\mathcal{F}_t^{W, \tau}$ -measurable function ϕ on $\mathcal{C} \times [0, T]$. Finally, the last step used the easy identity

$$\mathbb{E} \left[g_T(B, \mu) h^+(W) | \mathcal{F}_t^{B, W, \mu} \right] = \mathbb{E} \left[g_T(B, \mu) | \mathcal{F}_t^{B, \mu} \right] \mathbb{E}[h^+(W)]. \quad \square$$

7 Proofs of Limit Theorems

This section is devoted to the proofs of Theorems 4.4 and 4.5. At this point it may be useful to recall the notations in these theorems. In particular, take care to distinguish the empirical measure of stopping times $\bar{\mu}(\tau^n)$, defined in (3.1), from the *joint* empirical measure $\hat{\mu}(\tau^n)$, defined in (4.3).

7.1 Proof of Theorem 4.4

Abbreviate $\hat{\mu}^n = \hat{\mu}^n(\tau^n)$. Note first that $\mathbb{P} \circ (B, W^k)^{-1} = \mathcal{W}^2$ for all k , so the $\mathcal{C}^2$ -marginal of P_n does not depend on n . Clearly, the $[0, T]$ -marginal sequence $(P_n \circ \tau^{-1})_{n=1}^\infty$ is tight because $[0, T]$ is compact. To show that the marginal sequence $(P_n \circ (\hat{\mu}^n)^{-1})_{n=1}^\infty$ is tight, it suffices to show that the sequence of mean measures $\mathbb{E}^{P_n}[\hat{\mu}^n(\cdot)] \in \mathcal{P}(\mathcal{W} \times [0, T])$ is tight (c.f. the proof of [34, Proposition 2.2]). But this follows from the observation that the first marginal of $\mathbb{E}^{P_n}[\hat{\mu}^n(\cdot)]$ is the Wiener measure for each n . As each of marginal sequences is tight, the sequence $(P_n)_{n=1}^\infty \subset \mathcal{P}(\Omega)$ is tight. Let P be any limit point of P_n , and relabel the subsequence to assume $P_n \rightarrow P$. We check that P satisfies the five defining properties of a weak MFE.

Proof of (1) First, note that

$$P \circ (B, W)^{-1} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{P} \circ (B, W^k)^{-1} = \mathcal{W}^2.$$

We next prove that (B, W) is a Wiener process with respect to $\mathbb{F}_+^{B, W, \mu, \tau}$, or equivalently with respect to $\mathbb{F}^{B, W, \mu, \tau}$. Fix $t \in [0, T]$. Let f_t, g_t, h_t , and h^+ be bounded continuous functions on $[0, T]$, $\mathcal{P}(\mathcal{C} \times [0, T])$, $\mathcal{C}^2$, and $\mathcal{C}^2$, respectively. Assume f_t is $\mathcal{F}_t^\tau$ -measurable, g_t is $\mathcal{F}_t^\mu$ -measurable, h_t is $\mathcal{F}_t^{B, W}$ -measurable, and h^+ is $\sigma\{(B_s - B_t, W_s - W_t) : s \in [t, T]\}$ -measurable. Then, because $(B, W^1, \dots, W^n)$ are

$\mathbb{F}_+^{B, W^1, \dots, W^n}$ -Wiener processes,

$$\begin{aligned} & \mathbb{E}^P [f_t(\tau) g_t(\mu) h_t(B, W) h^+(B, W)] \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} [f_t(\tau_i^n) g_t(\widehat{\mu}^n) h_t(B, W^i) h^+(B, W^i)] \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} [f_t(\tau_i^n) g_t(\widehat{\mu}^n) h_t(B, W^i)] \mathbb{E}^{\mathbb{P}} [h^+(B, W^i)] \\ &= \mathbb{E}^P [f_t(\tau) g_t(\mu) h_t(B, W)] \mathbb{E}^P [h^+(B, W)]. \end{aligned}$$

This shows that $\sigma\{(B_s - B_t, W_s - W_t) : s \in [t, T]\}$ is independent of $\mathcal{F}_t^{B, W, \mu, \tau}$, under P .

Proof of (2) To show that (B, μ) and W are independent under P is straightforward: For bounded continuous functions $f : \mathcal{C} \times \mathcal{P}(\mathcal{C} \times [0, T]) \rightarrow \mathbb{R}$ and $g : \mathcal{C} \rightarrow \mathbb{R}$, the law of large numbers yields

$$\begin{aligned} & \mathbb{E}^P [f(B, \mu) g(W)] - \mathbb{E}^P [f(B, \mu)] \mathbb{E}^P [g(W)] \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} [f(B, \widehat{\mu}^n) g(W^i)] - \mathbb{E}^{\mathbb{P}} [f(B, \widehat{\mu}^n)] \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} [g(W^i)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}} \left[f(B, \widehat{\mu}^n) \left(\frac{1}{n} \sum_{i=1}^n g(W^i) - \int g d\mathcal{W} \right) \right] \\ &= 0, \end{aligned}$$

since W^i are i.i.d. with law $\mathcal{W}$ under $\mathbb{P}$.

Proof of (5) The proof of the fixed point condition (5) is also straightforward. Let f and g be bounded continuous functions on $\mathcal{C} \times \mathcal{P}(\mathcal{C} \times [0, T])$ and $\mathcal{C} \times [0, T]$, respectively, and notice that

$$\begin{aligned} \mathbb{E}^P [f(B, \mu) g(W, \tau)] &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} [f(B, \widehat{\mu}^n) g(W^i, \tau_i^n)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}} \left[f(B, \widehat{\mu}^n) \int g d\widehat{\mu}^n \right] \\ &= \mathbb{E}^P \left[f(B, \mu) \int g d\mu \right]. \end{aligned}$$

Proof of (3) Because we have established properties (2) and (5), Lemma 6.7 will immediately yield (3) once we can show that $\mathcal{F}_T^{B, \mu} \vee \mathcal{F}_t^{W, \tau}$ is independent of $\sigma\{W_s - W_t : s \in [t, T]\}$, for every $t \in [0, T]$. Fix $t \in [0, T]$. Fix bounded continuous functions f on $\mathcal{C} \times \mathcal{P}(\mathcal{C} \times [0, T])$, g_t on $[0, T]$, h_t on $\mathcal{C}$, and h^+ on $\mathcal{C}$. Assume f is uniformly continuous (thus $\mathcal{F}_T^{B, \mu}$ -measurable), g is $\mathcal{F}_t^{\tau}$ -measurable, h_t is $\mathcal{F}_t^W$ -measurable, and

h^+ is $\sigma\{W_s - W_t : s \in [t, T]\}$ -measurable. Define

$$\widehat{\mu}^{n,-i} := \frac{1}{n-1} \sum_{k \neq i} \delta_{(W^k, \tau_k^n)} = \frac{n}{n-1} \widehat{\mu}^n - \frac{1}{n-1} \delta_{(W^k, \tau_k^n)},$$

and note that $\|\widehat{\mu}^{n,-i} - \widehat{\mu}^n\|_{TV} \leq 2/(n-1)$ a.s. The total variation topology is finer than weak convergence, and so

$$|f(B, \widehat{\mu}^n) - f(B, \widehat{\mu}^{n,-i})| \rightarrow 0,$$

in L^∞ , uniformly in i . Now, since $\sigma\{W_s^i - W_t^i : s \in [t, T]\}$ is independent of $\mathcal{F}_t^{W^i} \vee \mathcal{F}_T^{B, (W^k)_{k \neq i}}$, we have

$$\begin{aligned} & \mathbb{E}^P [f(B, \mu) g_t(\tau) h_t(W) h^+(W)] \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^P [f(B, \widehat{\mu}^n) g_t(\tau_i^n) h_t(W^i) h^+(W^i)] \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^P [f(B, \widehat{\mu}^{n,-i}) g_t(\tau_i^n) h_t(W^i) h^+(W^i)] \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^P [f(B, \widehat{\mu}^{n,-i}) g_t(\tau_i^n) h_t(W^i)] \mathbb{E}^P [h^+(W^i)] \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^P [f(B, \widehat{\mu}^n) g_t(\tau_i^n) h_t(W^i)] \mathbb{E}^P [h^+(W^i)] \\ &= \mathbb{E}^P [f(B, \mu) g_t(\tau) h_t(W)] \mathbb{E}^P [h^+(W)]. \end{aligned}$$

This implies $\mathcal{F}_T^{B, \mu} \vee \mathcal{F}_t^{W, \tau}$ is independent of $\sigma\{W_s - W_t : s \in [t, T]\}$, under P .

Proof of (4) It remains to prove that the optimality condition holds in the limit. Recall that $\mu^\tau(\cdot) = \mu(\mathcal{C} \times \cdot)$ denotes the $[0, T]$ -marginal of μ . By Corollary 6.5, it suffices to show

$$\mathbb{E}^P [F(B, W, \mu^\tau, \tau)] \geq \mathbb{E}^P [F(B, W, \mu^\tau, \sigma)]$$

for every $\mathbb{F}^{B, W, \mu}$ -stopping time σ on Ω of the form $\sigma = \hat{\sigma}(B, W, \mu)$, where $\hat{\sigma} : \mathcal{C}^2 \times \mathcal{P}(\mathcal{C} \times [0, T]) \rightarrow [0, T]$ is continuous. Fix such a stopping time. For the n -player game define

$$\sigma_i = \hat{\sigma}(B, W^i, \widehat{\mu}^n(\tau^n)).$$

Then σ_i is a $\mathbb{F}^{B, W^1, \dots, W^n}$ -stopping time. Recall that $\bar{\mu}^n$ denotes the $[0, T]$ -marginal of the joint empirical measure $\hat{\mu}^n$. The Nash property implies

$$\begin{aligned} \mathbb{E}^P[F(B, W, \mu, \tau)] &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} \left[F(B, W^i, \bar{\mu}^n(\tau^n), \tau_i^n) \right] \\ &\geq \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} \left[F(B, W^i, \bar{\mu}^n(\tau^{n,-i}, \sigma_i), \sigma_i) \right] \\ &= \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} \left[F(B, W^i, \bar{\mu}^n(\tau^n), \hat{\sigma}(B, W^i, \hat{\mu}^n(\tau^n))) \right] \\ &= \mathbb{E}^P[F(B, W, \mu^\tau, \hat{\sigma}(B, W, \mu))]. \end{aligned}$$

Indeed, the equality in the third line follows from the easy estimate $\|\bar{\mu}^n(\tau^{n,-i}, \sigma_i) - \bar{\mu}^n(\tau^n)\|_{TV} \leq 2/n$, where $\|m\|_{TV} = \sup_{|f| \leq 1} \int f dm$ denotes total variation, and also from the continuity of $F = F(b, w, m, t)$ in m ensured by assumption **D**. Both the first and last lines use Lemma 6.1 to deal with the potential discontinuity of F in (B, W) , with the last step using crucially the continuity of $\hat{\sigma}$. $\square$

7.2 Proof of Theorem 4.5

Let $P \in \mathcal{P}(\Omega)$ be a weak MFE. Construct, on some alternative probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$, a $\mathcal{C} \times \mathcal{P}(\mathcal{C} \times [0, T])$ -valued random variable (B, μ) with law $P \circ (B, \mu)^{-1}$ and a sequence of random variables (W^i, τ^i) , which are conditionally independent given (B, μ) and have common conditional law μ . There is some abuse of notation here, as (B, μ) is used both for the new random variable and for the random variable defined on the canonical space Ω , but this should cause no confusion as we work exclusively on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ in this proof.

The law of (B, μ, W^i, τ^i) is precisely P , for each i . As usual, let $\tau^n = (\tau^1, \dots, \tau^n)$, and for $t^1, \dots, t^n \in [0, T]$ define the empirical measures (now on $\tilde{\Omega}$)

$$\bar{\mu}^n(t^1, \dots, t^n) = \frac{1}{n} \sum_{i=1}^n \delta_{t^i}, \quad \hat{\mu}^n(t^1, \dots, t^n) = \frac{1}{n} \sum_{i=1}^n \delta_{(W^i, t^i)}.$$

Define

$$\epsilon_n := \left\{ \sup_{\sigma \in \mathcal{S}_n} \mathbb{E} \left[F(B, W^1, \bar{\mu}^n(\tau^{n,-1}, \sigma), \sigma) \right] - \mathbb{E} \left[F(B, W^1, \bar{\mu}^n(\tau^n), \tau^1) \right] \right\}^+,$$

where $\mathcal{S}_n$ is the set of $\mathbb{F}^{B, W^1, \dots, W^n}$ -stopping times (defined on $\tilde{\Omega}$). By symmetry

$$\sup_{\sigma \in \mathcal{S}_n} \mathbb{E} \left[F(B, W^k, \bar{\mu}^n(\tau^{n,-k}, \sigma), \sigma) \right] \leq \epsilon_n + \mathbb{E} \left[F(B, W^k, \bar{\mu}^n(\tau^n), \tau^k) \right], \quad (7.1)$$

for every $k = 1, \dots, n$. We first show that $\epsilon_n \rightarrow 0$. Indeed,

$$\epsilon_n \leq \mathbb{E} \left[\sup_{t \in [0, T]} \left| F(B, W^1, \bar{\mu}^n(\tau^{n,-1}, t), t) - F(B, W^1, \mu, t) \right| \right] + \left\{ \sup_{\sigma \in \mathcal{S}_\infty} \mathbb{E} \left[F(B, W^1, \mu, \sigma) \right] - \mathbb{E} \left[F(B, W^1, \bar{\mu}^n(\tau^n), \tau^1) \right] \right\}^+,$$

where $\mathcal{S}_\infty = \cup_{n \geq 1} \mathcal{S}_n$. We claim that the term on the first line converges to zero. Indeed, $\|\bar{\mu}^n(\tau^{n,-1}, t) - \bar{\mu}^n(\tau^n)\| \leq 2/n$, and $\bar{\mu}^n(\tau^n) \rightarrow \mu$ weakly a.s. by the (conditional) law of large numbers. Conclude using the assumption that $\mathcal{P}([0, T]) \times [0, T] \ni (m, t) \mapsto F(b, w, m, t)$ is (uniformly) continuous for each fixed (b, w) . The second term also tends to zero, because

$$\mathbb{E} \left[F(B, W^1, \bar{\mu}^n(\tau^n), \tau^1) \right] \rightarrow \mathbb{E} \left[F(B, W^1, \mu, \tau^1) \right] \geq \sup_{\sigma \in \mathcal{S}_\infty} \mathbb{E} \left[F(B, W^1, \mu, \sigma) \right].$$

Indeed, to prove the last inequality, note that for any $\sigma \in \mathcal{S}_\infty$ we can easily check that $\mathcal{F}_t^\sigma$ is conditionally independent of $\mathcal{F}_T^{B, W^1, \mu}$ given $\mathcal{F}_t^{B, W^1, \mu}$, for every t . Because (B, W^1, μ, τ^1) has law P , which is a weak MFE, the optimality condition (4) of Definition 4.2 provides the desired inequality.

It may appear that we have shown that $(\tau^1, \dots, \tau^n)$ form an ϵ_n -Nash equilibrium for the n -player game, with $\epsilon_n \rightarrow 0$, but this is not accurate. The stopping times τ^i are not stopping times relative to $\mathbb{F}^{B, W^1, \dots, W^n}$, but rather to a larger filtration. This necessitates one more approximation, using a straightforward extension of Theorem 6.4 to deal with vectors of stopping times as opposed to single stopping times; the proof of this extension is exactly the same but notationally more cumbersome. Note that $\mathcal{F}_t^{\tau^1, \dots, \tau^n}$ is conditionally independent of $\mathcal{F}_T^{B, W^1, \dots, W^n}$ given $\mathcal{F}_t^{B, W^1, \dots, W^n, \mu}$, for every t , simply because $(B, W^1, \dots, W^n)$ is a $\mathbb{F}^{B, W^1, \dots, W^n, \tau^1, \dots, \tau^n}$ -Wiener process. Hence, using the aforementioned extension of Theorem 6.4, we may find $\mathbb{F}^{B, W^1, \dots, W^n}$ -stopping times, $\tau_k^1, \dots, \tau_k^n$, such that

$$(B, W^1, \dots, W^n, \tau_k^1, \dots, \tau_k^n) \Rightarrow (B, W^1, \dots, W^n, \tau^1, \dots, \tau^n),$$

as $k \rightarrow \infty$. Let $\tau^{n,k} = (\tau_k^1, \dots, \tau_k^n)$, and define

$$\epsilon_n^k := \max_{i=1, \dots, n} \left\{ \sup_{\sigma \in \mathcal{S}_n} \mathbb{E} \left[F(B, W^i, \bar{\mu}^n(\tau^{n,k,-i}, \sigma), \sigma) \right] - \mathbb{E} \left[F(B, W^i, \bar{\mu}^n(\tau^{n,k}), \tau_k^i) \right] \right\}^+.$$

For a fixed n , we can argue that $\lim_{k \rightarrow \infty} \epsilon_n^k = \epsilon_n$. Indeed, this follows from the observation that

$$\lim_{k \rightarrow \infty} \max_{i=1, \dots, n} \left| \mathbb{E} \left[F(B, W^i, \bar{\mu}^n(\tau^{n,k}), \tau_k^i) \right] - \mathbb{E} \left[F(B, W^i, \bar{\mu}^n(\tau^n), \tau^i) \right] \right| = 0,$$

by construction, and

$$\lim_{k \rightarrow \infty} \max_{i=1, \dots, n} \mathbb{E} \times \left[\sup_{t \in [0, T]} \left| F(B, W^i, \bar{\mu}^n(\tau^{n, k, -i}, t), t) - F(B, W^i, \bar{\mu}^n(\tau^{n, -i}, t), t) \right| \right] = 0,$$

because $m \mapsto F(b, w, m, t)$ is continuous in m , *uniformly in t* , for almost every fixed (b, w) .

In summary, we may find $k_n \rightarrow \infty$ such that

$$\frac{1}{n} \sum_{i=1}^n \text{Law}(B, W^i, \hat{\mu}^n(\tau^{n, k_n}), \tau_{k_n}^i) \rightarrow \text{Law}(B, W^1, \mu, \tau^1),$$

as well as $\epsilon_n^{k_n} \downarrow 0$ and

$$\begin{aligned} & \max_{i=1, \dots, n} \sup_{\sigma \in S_n} \mathbb{E} \left[F(B, W^i, \bar{\mu}^n(\tau^{n, k_n, -i}, \sigma), \sigma) \right] \\ & \leq \mathbb{E} \left[F(B, W^i, \bar{\mu}^n(\tau^{n, k_n}), \tau_{k_n}^i) \right] + \epsilon_n^{k_n}. \end{aligned}$$

8 Existence Under Continuity Assumption D

This section is devoted to the proof of Theorem 4.6. First, we prove existence of a weak MFE (more precisely, a strong MFE with weak stopping time) in the case that the time set and the common noise range space are finite. Then, we take weak limits. We introduce the following discretization.

For each positive integer n , let $t_i^n = i2^{-n}T$ for $i = 0, \dots, 2^n$. Choose partitions π^n of $\mathbb{R}$ into n measurable sets of strictly positive Lebesgue measure, in such a way that π^{n+1} is a refinement of π^n for each n and the union $\bigcup_{n \geq 1} \pi^n$ generates the entire Borel σ -field. We will define a sub-filtration of $\mathbb{F}^B$ according to which sets of π^n contains the increments $B_{t_{i+1}^n} - B_{t_i^n}$. Precisely, define the σ -field

$$\begin{aligned} \mathcal{G}_{t_k^n}^n &= \sigma \left\{ \{B_{t_i^n} - B_{t_{i-1}^n} \in C\} : C \in \pi^n, i = 1, \dots, k \right\}, \text{ for } k = 1, \dots, 2^n, \\ \mathcal{G}_0^n &= \sigma \{ \{B_0 \in C\} : C \in \pi^n \}. \end{aligned}$$

Additionally, for $t \in [0, T]$, define $\lfloor t \rfloor_n = \max\{t_k^n : k = 0, \dots, 2^n, t_k^n \leq t\}$, and let $\mathcal{G}_t^n = \mathcal{G}_{\lfloor t \rfloor_n}^n$. Then $\mathbb{G}^n = (\mathcal{G}_t^n)_{t \in [0, T]}$ defines a filtration, and $\mathcal{G}_t^n$ is finite for each n, t . Moreover,

$$\mathcal{F}_t^B = \sigma \left(\bigcup_{n=1}^{\infty} \mathcal{G}_t^n \right), \text{ for } t \in [0, T].$$

By construction, $\mathcal{W}(C) > 0$ for every nonempty set $C \in \mathcal{G}_T^n$.

Let $\mathcal{M}_n$ denote the set of functions $\bar{m} : \mathcal{C} \rightarrow \mathcal{P}([0, T])$ such that, for each $t \in [0, T]$ the map $b \mapsto \bar{m}(b)[0, t]$ is $\mathcal{G}_t^n$ -measurable. Of course, since $\mathcal{G}_T^n$ is finite, a map $\bar{m} \in \mathcal{M}_n$ must be constant on each atom of $\mathcal{G}_T^n$. Endowed with the topology of pointwise convergence, $\mathcal{M}_n$ is easily seen to be compact, because $\mathcal{P}([0, T])$ is compact. Lastly, define $\mathcal{A}$ as the set of probability measures $Q \in \mathcal{P}(\mathcal{C} \times [0, T])$ under which B and W are independent $\mathbb{F}^{B, W, \tau}$ -Wiener processes.

Theorem 8.1 *For each n , there exist $\bar{m} \in \mathcal{M}_n$ and $Q \in \mathcal{A}$ satisfying the following:*

- (1) $\bar{m}(B) = Q(\tau \in \cdot | \mathcal{G}_T^n)$.
- (2) *The optimality condition holds,*

$$\mathbb{E}^Q[F(B, W, \bar{m}(B), \tau)] \geq \sup_{Q' \in \mathcal{A}} \mathbb{E}^{Q'}[F(B, W, \bar{m}(B), \tau)]$$

Proof It follows from Theorem 6.4 that $\mathcal{A}$ is compact and convex; see Remark 6.3. Define a map Φ_n from $\mathcal{M}_n$ to subsets of $\mathcal{A}$ by

$$\Phi_n(\bar{m}) = \arg \max_{Q \in \mathcal{A}} \mathbb{E}^Q[F(B, W, \bar{m}(B), \tau)].$$

The map $(\bar{m}, Q) \mapsto \mathbb{E}^Q[F(B, W, \bar{m}(B), \tau)]$ is jointly continuous on $\mathcal{M}_n \times \mathcal{A}$, thanks to Lemma 6.1 and continuity of $F = F(b, w, m, t)$ in (m, t) . Thus, by Berge's theorem [4, Theorem 17.31], Φ_n has closed graph and takes nonempty convex values. Define a map $\Psi_n : \mathcal{M}_n \rightarrow 2^{\mathcal{M}_n}$ by

$$\Psi_n(\bar{m}) = \{Q(\tau \in \cdot | \mathcal{G}_T^n) : Q \in \Phi_n(\bar{m})\}.$$

Because $Q \circ B^{-1} = \mathcal{W}$, for continuous bounded functions $f : [0, T] \rightarrow \mathbb{R}$ we may write

$$\mathbb{E}^Q[f(\tau) | \mathcal{G}_T^n](b) = \sum_C \frac{\mathbb{E}^Q[f(\tau) 1_{\{B \in C\}}]}{\mathcal{W}(C)} 1_C(b), \text{ for } b \in \mathcal{C},$$

where the sum is over atoms of $\mathcal{G}_T^n$. For each such atom C , note that $\mathcal{W}(C) > 0$, and the map $Q \mapsto \mathbb{E}^Q[f(\tau) 1_{\{B \in C\}}]$ is continuous on $\mathcal{A}$ by Lemma 6.1. We may thus view $Q \mapsto Q(\tau \in \cdot | \mathcal{G}_T^n)$ as a continuous affine map from $\mathcal{A}$ to $\mathcal{M}_n$. Hence, the set-valued map Ψ_n has closed graph, and its values are nonempty, convex, and compact. By Kakutani's theorem [4, Corollary 17.55], it admits a fixed point; that is, there exists $\bar{m} \in \mathcal{M}_n$ such that $\bar{m} \in \Psi_n(\bar{m})$. $\square$

Proof of Theorem 4.6

For each n , let $\bar{m}_n \in \mathcal{M}_n$ and $Q_n \in \mathcal{A}$ satisfy properties (1–2) of Theorem 8.1. Define $\hat{m}_n : \mathcal{C} \rightarrow \mathcal{P}(\mathcal{C} \times [0, T])$ by

$$\hat{m}_n(B) = Q_n((W, \tau) \in \cdot | \mathcal{G}_T^n).$$

We note here that if $C \in \mathcal{F}_t^{W, \tau}$ for $t \in [0, T]$, then we have

$$\widehat{m}_n(B) = Q_n((W, \tau) \in \cdot | \mathcal{G}_{\lceil t \rceil_n}^n), \quad (8.1)$$

where $\lceil t \rceil_n = \min\{t_k^n : k = 0, \dots, 2^n, t_k^n \geq t\}$. Indeed, this holds because B is an $\mathbb{F}^{B, W, \tau}$ -Wiener process under Q_n , and because $\mathcal{G}_T^n$ is generated by $\mathcal{G}_{\lceil t \rceil_n}^n$ and the events $\{B_{t_i^n} - B_{t_{i-1}^n} \in C'\}$ for $C' \in \pi^n$ and for $t_{i-1}^n \geq t$.

Next, define $P_n \in \mathcal{P}(\Omega)$ by

$$P_n = Q_n \circ (B, W, \widehat{m}_n(B), \tau)^{-1}.$$

We claim that $(P_n)_{n=1}^\infty$ is tight and that every limit point is a weak MFE in the sense of Definition 4.2. We begin with tightness, by showing that each marginal is tight. Note that the first marginal $P_n \circ (B, W)^{-1} = \mathcal{W}^2$ is clearly tight, as it is constant. Moreover, $P_n \circ \tau^{-1}$ is tight because $[0, T]$ is compact. To prove that $P_n \circ \mu^{-1}$ is tight, it suffices to show that the mean measures $\mathbb{E}^{P_n}[\mu(\cdot)]$ are tight (c.f. the proof of [34, Proposition 2.2]). But, for any measurable set $C \subset \mathcal{C} \times [0, T]$,

$$\begin{aligned} \mathbb{E}^{P_n}[\mu(C)] &= \mathbb{E}^{Q_n}[\widehat{m}_n(B)(C)] = \mathbb{E}^{Q_n}[Q_n((W, \tau) \in C | \mathcal{G}_T^n)] \\ &= Q_n((W, \tau) \in C) = P_n((W, \tau) \in C). \end{aligned}$$

That is, the mean measure $\mathbb{E}^{P_n}[\mu(\cdot)]$ is precisely $P_n \circ (W, \tau)^{-1}$, which we already observed to be tight.

With tightness of $(P_n)_{n=1}^\infty$ established, fix a limit point P and abuse notation by assuming $P_n \rightarrow P$. We will show that P is a weak MFE by checking the five properties of Definition 4.2:

Proof of (1) Clearly $P \circ (B, W)^{-1} = \lim_{n \rightarrow \infty} Q_n \circ (B, W)^{-1} = \mathcal{W}^2$. We must show that (B, W) are $\mathbb{F}^{B, W, \mu, \tau}$ -Wiener processes under P . Fix N and fix $t \in \{t_0^N, \dots, t_{2^N-1}^N\}$. Let f_t, g_t, h_t , and h^+ be bounded continuous functions on $[0, T], \mathcal{P}(\mathcal{C} \times [0, T]), \mathcal{C}^2$, and $\mathcal{C}^2$, respectively. Assume f_t is $\mathcal{F}_t^\tau$ -measurable, g_t is $\mathcal{F}_t^\mu$ -measurable, h_t is $\mathcal{F}_t^{B, W}$ -measurable, and h^+ is $\sigma\{(B_s - B_t, W_s - W_t) : s \in [t, T]\}$ -measurable. Since $\lceil t \rceil_n = t$ for $n \geq N$ and $\mathcal{G}_t^n \subset \mathcal{F}_t^B$, it follows from (8.1) that the map $b \mapsto \widehat{m}_n(b)(C)$ is $\mathcal{F}_t^B$ -measurable for every $C \in \mathcal{F}_t^{W, \tau}$. This implies $\widehat{m}_n$ is $\mathcal{F}_t^B / \mathcal{F}_t^\mu$ -measurable and, in particular, $g_t(\widehat{m}_n(B))$ is $\mathcal{F}_t^B$ -measurable. Thus, since (B, W) is an $\mathbb{F}^{B, W, \tau}$ -Wiener process under Q_n ,

$$\begin{aligned} \mathbb{E}^P[f_t(\tau)g_t(\mu)h_t(B, W)h^+(B, W)] &= \lim_{n \rightarrow \infty} \mathbb{E}^{Q_n}[f_t(\tau)g_t(\widehat{m}_n(B))h_t(B, W)h^+(B, W)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{Q_n}[f_t(\tau)g_t(\widehat{m}_n(B))h_t(B, W)] \mathbb{E}^{Q_n}[h^+(B, W)] \\ &= \mathbb{E}^P[f_t(\tau)g_t(\mu)h_t(B, W)] \mathbb{E}^P[h^+(B, W)]. \end{aligned}$$

This is enough to conclude that $\sigma\{(B_s - B_t, W_s - W_t) : s \in [t, T]\}$ is independent of $\mathcal{F}_t^{B, W, \mu, \tau}$, under P . We have only shown this to be true for $t \in \bigcup_{N=1}^{\infty} \{t_0^N, \dots, t_{2^N-1}^N\}$, but this it suffices to note that this set is dense in $[0, T]$.

Proof of (2) Because B and W are independent under Q_n , for bounded continuous functions f on $\mathcal{C} \times \mathcal{P}(\mathcal{C} \times [0, T])$ and g on $\mathcal{C}$ we have

$$\begin{aligned} \mathbb{E}^P[f(B, \mu)g(W)] &= \lim_{n \rightarrow \infty} \mathbb{E}^{Q_n}[f(B, \widehat{m}(B))g(W)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{Q_n}[f(B, \widehat{m}(B))]\mathbb{E}^{Q_n}[g(W)] \\ &= \mathbb{E}^P[f(B, \mu)]\mathbb{E}^P[g(W)]. \end{aligned}$$

Thus (B, μ) and W are independent under P .

Proof of (5) Fix $N \geq 1$, and let f be a bounded $\mathcal{G}_T^N$ -measurable function on $\mathcal{C}$. Let g and h be continuous bounded functions on $\mathcal{C} \times [0, T]$ and $\mathcal{P}(\mathcal{C} \times [0, T])$, respectively. Noting that f is $\mathcal{G}_T^n$ -measurable for all $n \geq N$, and using Lemma 6.1 to deal with the discontinuity of f in b , we have

$$\begin{aligned} \mathbb{E}^P[f(B)h(\mu)g(W, \tau)] &= \lim_{n \rightarrow \infty} \mathbb{E}^{Q_n}[f(B)h(\widehat{m}_n(B))g(W, \tau)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{Q_n}\left[f(B)h(\widehat{m}_n(B)) \int g d\widehat{m}_n(B)\right] \\ &= \mathbb{E}^P\left[f(B)h(\mu) \int g d\mu\right]. \end{aligned}$$

This holds for each N and each $\mathcal{G}_T^N$ -measurable f . Since $\mathcal{F}_T^B = \sigma(\bigcup_{n \geq 1} \mathcal{G}_T^n)$, the same identity must hold for every $\mathcal{F}_T^B$ -measurable f .

Proof of (3) Because we have established properties (2) and (5), Lemma 6.7 will yield (3) once we can show that $\mathcal{F}_T^{B, \mu} \vee \mathcal{F}_t^{W, \tau}$ is independent of $\sigma\{W_s - W_t : s \in [t, T]\}$, for every $t \in [0, T]$. Fix $t \in [0, T]$. Fix bounded continuous functions f on $\mathcal{C} \times \mathcal{P}(\mathcal{C} \times [0, T])$, g_t on $[0, T]$, h_t on $\mathcal{C}$, and h^+ on $\mathcal{C}$. Assume g is $\mathcal{F}_t^{\tau}$ -measurable, h_t is $\mathcal{F}_t^W$ -measurable, and h^+ is $\sigma\{W_s - W_t : s \in [t, T]\}$ -measurable. Under Q_n , (B, W) is a standard $\mathbb{R}^{B, W, \tau}$ -Wiener process, and it follows easily that $h^+(W)$ is independent of $\mathcal{F}_T^B \vee \mathcal{F}_t^{W, \tau}$. Thus,

$$\begin{aligned} &\mathbb{E}^P[f(B, \mu)g_t(\tau)h_t(W)h^+(W)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{Q_n}[f(B, \widehat{m}_n(B))g_t(\tau)h_t(W)h^+(W)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{Q_n}[f(B, \widehat{m}_n(B))g_t(\tau)h_t(W)]\mathbb{E}^{Q_n}[h^+(W)] \\ &= \mathbb{E}^P[f(B, \mu)g_t(\tau)h_t(W)]\mathbb{E}^P[h^+(W)]. \end{aligned}$$

This shows $\mathcal{F}_T^{B, \mu} \vee \mathcal{F}_t^{W, \tau}$ is independent of $\sigma\{W_s - W_t : s \in [t, T]\}$, under P .

Proof of (4) By Corollary 6.5, it suffices to show

$$\mathbb{E}^P[F(B, W, \mu^{\tau}, \tau)] \geq \mathbb{E}^P[F(B, W, \mu^{\tau}, \sigma)]$$

for every $\mathbb{F}^{B,W,\mu}$ -stopping time σ of the form $\sigma = \hat{\sigma}(B, W, \mu)$, where $\hat{\sigma} : \mathcal{C}^2 \times \mathcal{P}(\mathcal{C} \times [0, T]) \rightarrow [0, T]$ is continuous. Fix such a stopping time. Using the continuity of $\hat{\sigma}$, and using Lemma 6.1 to handle the discontinuity of F in (B, W) , we have

$$\begin{aligned} & \mathbb{E}^P[F(B, W, \mu^\tau, \tau)] - \mathbb{E}^P[F(B, W, \mu^\tau, \sigma)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{Q_n}[F(B, W, \bar{m}_n(B), \tau)] - \mathbb{E}^{Q_n}[F(B, W, \bar{m}_n(B), \sigma(B, W, \hat{m}_n(B)))] \\ &\geq 0, \end{aligned}$$

where we have finally used the optimality property of $\bar{m}_n$ from Theorem 8.1. Indeed, the law

$$Q'_n := Q_n \circ (B, W, \sigma(B, W, \hat{m}_n(B)))^{-1}$$

is easily seen to belong to $\mathcal{A}$, by noting that (as we saw in the proof of (1)) $\hat{m}_n$ is adapted in the sense that $b \mapsto \hat{m}_n(b)(C)$ is $\mathcal{F}_t^B$ -measurable for $C \in \mathcal{F}_t^{W,\tau}$. $\square$

Acknowledgements We would like to thank Geoffrey Zhu for enlightening discussions at an early stage of our investigation of mean field games of timing.

Appendix A: Proof of Proposition 6.2

To prove Theorem 6.2, we need a preliminary results, borrowed from previous works of the authors. Recall that $\Rightarrow$ denotes convergence in law.

Proposition A.1 (Proposition C.1 of [12]) *Let X and Y be random variables defined on a common probability space, taking values in some Polish spaces E and F . If the law of X is nonatomic, and if F is (homeomorphic to) a convex subset of a locally convex space, then there exists a sequence of continuous functions $\phi_n : E \rightarrow F$ such that $(X, \phi_n(X)) \Rightarrow (X, Y)$.*

Proposition 6.2 extends Proposition A.1 to the dynamic setting, and this is where the role of compatibility is the clearest. This is contained in the third author's PhD thesis [27, Proposition 2.1.6], which itself was implicitly present in the proof of [12, Lemma 3.11], though we include the proof for the sake of completeness.

Proof of Proposition 6.2

The proof is an inductive application of Proposition A.1. First, in light of the assumption that the law of Y_1 is nonatomic, Proposition A.1 allows us to find a sequence of continuous functions $h_1^j : \mathcal{Y} \rightarrow \mathcal{X}$ such that $(Y_1, h_1^j(Y_1)) \Rightarrow (Y_1, X_1)$ as $j \rightarrow \infty$. Let us show that in fact $(Z, h_1^j(Y_1))$ converges to (Z, X_1) . Let $\phi : \mathcal{Z} \rightarrow \mathbb{R}$ be bounded and measurable, and let $\psi : \mathcal{X} \rightarrow \mathbb{R}$ be continuous. Note that Z and X_1 are conditionally independent given Y_1 , by assumption. Now use Lemma 6.1 to get

$$\begin{aligned}
 \lim_{j \rightarrow \infty} \mathbb{E}[\phi(Z)\psi(h_1^j(Y_1))] &= \lim_{j \rightarrow \infty} \mathbb{E}\left[\mathbb{E}[\phi(Z)|Y_1]\psi(h_1^j(Y_1))\right] \\
 &= \mathbb{E}[\mathbb{E}[\phi(Z)|Y_1]\psi(X_1)] \\
 &= \mathbb{E}[\mathbb{E}[\phi(Z)|Y_1]\mathbb{E}[\psi(X_1)|Y_1]] \\
 &= \mathbb{E}[\mathbb{E}[\phi(Z)\psi(X_1)|Y_1]] \\
 &= \mathbb{E}[\phi(Z)\psi(X_1)].
 \end{aligned}$$

The class of functions of the form $\mathcal{Z} \times \mathcal{X} \ni (z, x) \mapsto \phi(z)\psi(x)$, where ϕ and ψ are as above, is convergence determining (see, e.g., [16, Proposition 3.4.6(b)]), and we conclude that $(Z, h_1^j(Y_1)) \Rightarrow (Z, X_1)$.

We proceed inductively as follows. Abbreviate $Y^n := (Z_1, \dots, Z_n)$ for each $n = 1, \dots, N$, noting $Y^N = Y$, and similarly define X^n . Suppose we are given $1 \leq n < N$ and continuous functions $g_k^j : \mathcal{Y}^k \rightarrow \mathcal{X}$, for $k \in \{1, \dots, n\}$ and $j \geq 1$, satisfying

$$\lim_{j \rightarrow \infty} (Z, g_1^j(Y^1), \dots, g_n^j(Y^n)) = (Z, X^n), \quad (\text{A.1})$$

where convergence is in distribution, as usual. We will show that there exist continuous functions $h_k^i : \mathcal{Y}^k \rightarrow \mathcal{X}$ for each $k \in \{1, \dots, n+1\}$ and $i \geq 1$ such that

$$\lim_{i \rightarrow \infty} (Z, h_1^i(Y^1), \dots, h_{n+1}^i(Y^{n+1})) = (Z, X_1, \dots, X_{n+1}). \quad (\text{A.2})$$

By Proposition A.1 there exists a sequence of continuous functions $\hat{g}^j : (\mathcal{Y}^{n+1} \times \mathcal{X}^n) \rightarrow \mathcal{X}$ such that

$$\lim_{j \rightarrow \infty} (Y^{n+1}, X^n, \hat{g}^j(Y^{n+1}, X^n)) = (Y^{n+1}, X^n, X_{n+1}) = (Y^{n+1}, X^{n+1}). \quad (\text{A.3})$$

We claim now that

$$\lim_{j \rightarrow \infty} (Z, X^n, \hat{g}^j(Y^{n+1}, X^n)) = (Z, X^n, X_{n+1}). \quad (\text{A.4})$$

Indeed, let ϕ , ψ_n , and ψ be bounded measurable functions on $\mathcal{Z}$, $\mathcal{X}^n$, and $\mathcal{X}$, respectively, with ψ_n and ψ continuous. Use the conditional independence of Z and (Y^{n+1}, X^{n+1}) given Y^{n+1} along with (A.3) and Lemma 6.1 to get

$$\begin{aligned}
 &\lim_{j \rightarrow \infty} \mathbb{E}[\phi(Z)\psi_n(X^n)\psi(\hat{g}^j(Y^{n+1}, X^n))] \\
 &= \lim_{j \rightarrow \infty} \mathbb{E}\left[\mathbb{E}\left[\phi(Z)|Y^{n+1}\right]\psi_n(X^n)\psi(\hat{g}^j(Y^{n+1}, X^n))\right] \\
 &= \mathbb{E}\left[\mathbb{E}\left[\phi(Z)|Y^{n+1}\right]\psi_n(X^n)\psi(X_{n+1})\right] \\
 &= \mathbb{E}\left[\mathbb{E}\left[\phi(Z)|Y^{n+1}\right]\mathbb{E}\left[\psi_n(X^n)\psi(X_{n+1})|Y^{n+1}\right]\right]
 \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E} \left[\mathbb{E} \left[\phi(Z) \psi_n(X^n) \psi(X_{n+1}) \mid Y^{n+1} \right] \right] \\
&= \mathbb{E} \left[\phi(Z) \psi_n(X^n) \psi(X_{n+1}) \right].
\end{aligned}$$

Again, the class of functions of the form $\mathcal{Z} \times \mathcal{X}^n \times \mathcal{X} \ni (z, x, x') \mapsto \phi(z) \psi_n(x) \psi(x')$, where ϕ , ψ_n , and ψ are as above, is convergence determining, and (A.4) follows.

By continuity of $\hat{g}^j$, the limit (A.1) implies that, for each j ,

$$\begin{aligned}
&\lim_{k \rightarrow \infty} (Z, g_1^k(Y^1), \dots, g_n^k(Y^n), \hat{g}^j(Y^{n+1}, g_1^k(Y^1), \dots, g_n^k(Y^n))) \\
&= (Z, X_1, \dots, X_n, \hat{g}^j(Y^{n+1}, X_1, \dots, X_n)) \\
&= (Z, X^n, \hat{g}^j(Y^{n+1}, X^n)).
\end{aligned} \tag{A.5}$$

Combining the two limits (A.4) and (A.5), we may find a subsequence j_k such that

$$\lim_{k \rightarrow \infty} (Z, g_1^{j_k}(Y^1), \dots, g_n^{j_k}(Y^n), \hat{g}^{j_k}(Y^{n+1}, g_1^{j_k}(Y^1), \dots, g_n^{j_k}(Y^n))) = (Z, X^n, X_{n+1}).$$

Define $h_\ell^k := h_\ell^{j_k}$ for $\ell = 1, \dots, n$ and $h_{n+1}^k(Y^{n+1}) := \hat{g}^{j_k}(Y^{n+1}, g_1^{j_k}(Y^1), \dots, g_n^{j_k}(Y^n))$ to complete the induction. $\square$

Appendix B: The Lattice of Stopping Times

In this section, we prove that the set $\mathcal{S}$ of stopping times defined in Sect. 5 is a complete lattice. Recall that $\mathcal{S}$ is defined as the set of (equivalence classes of a.s. equal) random times τ defined on the probability space $\Omega^{\text{com}} \times \Omega^{\text{ind}}$, which are stopping times with respect to the filtration $\mathbb{F}^{\text{sig}}$. Recall that the *essential supremum* of a family Φ of random variables is defined as the minimal (with respect to a.s. order) random variable exceeding a.s. each element of $\mathbb{T}$:

Theorem B.1 (Theorem A.33 of [17]) *Let Φ be a set of real-valued random variables. Then there exists a unique (up to a.s. equality) random variable $Z = \text{ess sup } \Phi$ such that $Z \geq X$ a.s. for each $X \in \Phi$, and also $Z \leq Y$ a.s. for every random variable Y satisfying $Y \geq X$ a.s. for every $X \in \Phi$. Moreover, there exists a countable set $\Phi_0 \subset \Phi$ such that $Z = \sup_{X \in \Phi_0} X$ a.s.*

Proof Existence and uniqueness is stated in [17, Theorem A.33], and the proof therein constructs the desired Φ_0 . $\square$

The essential infimum is defined analogously, or simply by $\text{ess inf } \Phi = -\text{ess sup } (-\Phi)$.

Theorem B.2 *The set $\mathcal{S}$ is a complete lattice.*

Proof Fix a set $\Phi \subset \mathcal{S}$. Define $Z = \text{ess sup } \Phi$ and find a countable set $\{\tau_n : n \geq 1\} \subset \Phi$ such that $Z = \sup_n \tau_n$ a.s. Define $\sigma_n = \max_{k=1, \dots, n} \tau_k$, so that σ_n is an increasing sequence of stopping times with $\sigma_n \uparrow Z$ a.s. The increasing limit of a sequence of stopping times is again a stopping time [13, Theorem IV.55(b)], so $Z \in \mathcal{S}$.

A similar argument applies to show that the essential infimum of Φ is also a stopping time, and the only difference is that this step crucially uses the right-continuity of the filtration $\mathbb{F}^{\text{sig}}$; indeed, while the supremum of a sequence of stopping times is always a stopping time, the infimum of a sequence of stopping times is only guaranteed to be a stopping time if the underlying filtration is right-continuous [13, Theorem IV.55(c)]. $\square$

References

1. Acemoglu, D., Jensen, M.K.: Robust comparative statics in large dynamic economies. Tech. report, National Bureau of Economic Research (2012)
2. Acemoglu, D., Jensen, M.K.: Aggregate comparative statics. *Games Econ. Behav.* **81**, 27–49 (2013)
3. Adlakha, S., Johari, R.: Mean field equilibrium in dynamic games with strategic complementarities. *Oper. Res.* **61**(4), 971–989 (2013)
4. Aliprantis, C., Border, K.: *Infinite Dimensional Analysis: A Hitchhiker's Guide*, 3rd edn. Springer, Berlin (2007)
5. Balbus, L., Dziwulski, P., Reffett, K., Wozny, L.P.: A qualitative theory of large games with strategic complementarities. SSRN 2208125 (2014)
6. Balbus, L., Reffett, K., Woźny, L.: Monotone equilibria in nonatomic supermodular games. *Comment Games Econ. Behav.* **94**, 182–187 (2015)
7. Baxter, J.R., Chacon, R.V.: Compactness of stopping times. *Probab. Theory Relat. Fields* **40**(3), 169–181 (1977)
8. Brémaud, P., Yor, M.: Changes of filtrations and of probability measures. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* **45**(4), 269–295 (1978)
9. Bryant, J.: A model of reserves, bank runs and deposit insurance. *J. Bank. Financ.* **4**, 335–344 (1980)
10. Carmona, R., Delarue, F.: Probabilistic analysis of mean-field games. *SIAM J. Control Optim.* **51**(4), 2705–2734 (2013)
11. Carmona, R., Delarue, E.: *Probabilistic Theory of Mean Field Games*. Springer, Berlin (2016)
12. Carmona, R., Delarue, F., Lacker, D.: Mean field games with common noise. *Ann. Probab.* **44**(6), 3740–3803 (2016)
13. Dellacherie, C., Meyer, P.A.: *Probability and Potential*. North-Holland Mathematics Studies. Elsevier, Amsterdam (1979)
14. Diamond, D.W., Dybvig, P.H.: Bank runs, deposit insurance, and liquidity. *J. Polit. Econ.* **91**, 401–419 (1983)
15. El Karoui, N., Nguyen, D.H., Jeanblanc-Piqué, M.: Compactification methods in the control of degenerate diffusions: existence of an optimal control. *Stochastics* **20**(3), 169–219 (1987)
16. Ethier, S.N., Kurtz, T.G.: *Markov Processes: Characterization and Convergence*, vol. 282. Wiley, New York (2009)
17. Föllmer, H., Schied, A.: *Stochastic Finance: An introduction in Discrete Time*. De Gruyter Textbook. De Gruyter, Berlin (2011)
18. Green, E.J., Lin, P.: Implementing efficient allocations in a model of financial intermediation. *J. Econ. Theory* **109**(1), 1–23 (2003)
19. He, Z., Xiong, W.: Dynamic debt runs. *Rev. Financ. Stud.* **25**(6), 1799–1843 (2012)
20. Huang, C.-F., Li, L.: Continuous time stopping games with monotone reward structures. *Math. Oper. Res.* **15**, 496–507 (1990)
21. Huang, M., Caines, P.E., Malhamé, R.P.: Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.* **6**, 221–252 (2006)
22. Jacod, J., Mémin, J.: Sur un type de convergence intermédiaire entre la convergence en loi et la convergence en probabilité. *Séminaire de probabilités de Strasbourg* **15**, 529–546 (1981)
23. Jacod, J., Mémin, J.: *Weak and Strong Solutions of Stochastic Differential Equations: Existence and Stability*. Stochastic Integrals. Springer, Berlin (1981)
24. Jeanblanc, M., Le Cam, Y.: *Immersion Property and Credit Risk Modelling. Optimality and Risk-Modern Trends in Mathematical Finance*. Springer, New York (2009)

25. Kurtz, T.G.: Weak and strong solutions of general stochastic models. *Electron. Commun. Probab.* **19**(58), 1–16 (2014)
26. Lacker, D.: Mean field games via controlled martingale problems: existence of Markovian equilibria. *Stoch. Process. Appl.* **125**(7), 2856–2894 (2015)
27. Lacker, D.: Stochastic differential mean field game theory, Ph.D. thesis, (2015)
28. Lacker, D.: A general characterization of the mean field limit for stochastic differential games. *Probab. Theory Relat. Fields* **165**(3), 581–648 (2016)
29. Lasry, J.M., Lions, P.L.: Mean field games. *Jpn. J. Math.* **2**, 229–260 (2007)
30. Milgrom, P., Roberts, J.: Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* **58**, 1255–1277 (1990)
31. Morris, S., Shin, H.S.: Unique equilibrium in a model of self-fulfilling currency attacks. *Am. Econ. Rev.* **88**, 587–597 (1998)
32. Nutz, M.: A mean field game of optimal stopping. arXiv preprint [arXiv:1605.09112](https://arxiv.org/abs/1605.09112) (2016)
33. Rochet, J.C., Vives, X.: Coordination failures and the lender of last resort. *J. Eur. Econ. Assoc.* **2**, 1116–1148 (2004)
34. Sznitman, A.S.: Topics in propagation of chaos, Ecole d’été de Probabilités de Saint-Flour XIX, Lecture Notes in Mathematics, vol. 1464, pp. 165–251. Springer, Berlin (1989)
35. Topkis, D.M.: Minimizing a submodular function on a lattice. *Oper. Res.* **26**(2), 305–321 (1978)
36. Topkis, D.M.: Equilibrium points in nonzero-sum n -person submodular games. *SIAM J. Control Optim.* **17**(6), 773–787 (1979)
37. Vives, X.: Nash equilibrium with strategic complementarities. *J. Math. Econ.* **19**(3), 305–321 (1990)
38. Yang, J., Qi, X.: The nonatomic supermodular game. *Games Econ. Behav.* **82**, 609–620 (2013)