10

Dual Effect for Multi-Agent Stochastic Input-Output Systems

10.1 Introduction

In stochastic optimal control, a key issue is that "solutions" are searched for in terms of "closed-loop control laws" over available information and, as a consequence, a major potential difficulty is the fact that present control may affect future available information. This is known as the "dual effect" of control, and has been discussed in §1.1.3, §1.2.1, §1.3.2 and §4.2.3. Following [13], we will characterize the maximal set of closed-loop control laws containing open-loop laws and for which the information provided by observations closed with such a feedback remains fixed.

For this purpose, we consider in §10.2 the following variant of the Witsenhausen intrinsic model in §9.2. A *multi-agent stochastic input-output system*, in short MASIOS, is a multi-agent stochastic control system as in §9.2 where the information of an agent is described by an observation mapping (a signal), and where measurability is w.r.t. (complete) partition fields and not to σ-fields (see §3.3.2 and §3.4.2). In parallel to the discussion on the precedence and information-memory relations for multi-agent stochastic control systems in §9.4, we introduce their counterparts for MASIOS, as well as a typology of MASIOS.

The counterpart of a policy in the Witsenhausen intrinsic model in §9.2 is a control law, that is, a random variable defined on the universe Ω. An admissible control law for a focal agent is one that is measurable w.r.t. the agent closed-loop observation after control (of all agents). A collection of control laws (indexed by the set of agents) induces a partition of the universe Ω. No open-loop dual effect holds true when all constant control laws induce the same fixed partition. This is the object of §10.3.

Thanks to the typology of MASIOS introduced in §10.2, we characterize in §10.3.3 classes of control laws for which the induced partition coincides with this fixed partition. Therefore, if we restrict a stochastic optimization problem to such *no dual effect* control laws, the discretization of the control laws domain can be made in advance.

10.2 Multi-Agent Stochastic Input-Output Systems (MASIOS)

We introduce a *multi-agent stochastic input-output system*, which is a multi-agent stochastic control system as in §9.2, but where the information of an agent is described by an observation mapping (a signal), and where measurability is w.r.t. (complete) partition fields and not to σ-fields. We provide state models, especially linear ones, as examples inducing MASIOS. In parallel to the discussion on the precedence and information-memory relations for multi-agent stochastic control systems in §9.4, we introduce their counterparts for MASIOS, as well as a typology of MASIOS.

10.2.1 Multi-Agent Stochastic Input-Output Systems

Let A be a finite set representing agents. Each agent $\alpha \in A$ is supposed to make only one decision $u_\alpha \in \mathbb{U}_\alpha$, where $\mathbb{U}_\alpha$ is the control set for agent α, equipped with the complete π-field $\mathcal{U}_\alpha = 2^{\mathbb{U}_\alpha}$. Let Ω (*universe* or *sample space*) be a measurable set, with the complete π-field $\mathcal{F} = 2^{\Omega}$, which represents all uncertainties: any $\omega \in \Omega$ is called a *state of Nature*.

Remark 10.1. We adopt the same formalism as in Chapter 9, but for measurability which, here, is w.r.t. (complete) partition fields and not to σ-fields. We refer the reader to §3.3.2 for details. This option makes statements more compact and proofs more intuitive as compared to measurability w.r.t. σ-fields. $\lozenge$

As in (9.1) and (9.2), we define the *decision set* $\mathbb{U}_A$, and we equip it with the complete product π-field $\mathcal{U}_A$ (see Remark 3.12), called *decision field*:

$$\mathbb{U}_A := \prod_{\alpha \in A} \mathbb{U}_\alpha \; , \quad \mathcal{U}_A := \bigotimes_{\alpha \in A} \mathcal{U}_\alpha \; . \tag{10.1}$$

The *history space* $\mathbb{H}$ and its associated complete product π-field $\mathcal{H}$, called *history field*, are:

$$\mathbb{H} := \mathbb{U}_A \times \Omega \; , \quad \mathcal{H} := \mathcal{U}_A \otimes \mathcal{F} \; . \tag{10.2}$$

To each agent $\alpha \in A$ is attached an *observation function*

$$o_\alpha : \mathbb{H} \to \mathbb{Y}_\alpha \; . \tag{10.3}$$

Remark 10.2. Here, the information of agent α is described by a mapping $o_\alpha : \mathbb{H} \to \mathbb{Y}_\alpha$ defined over the history space $\mathbb{H} = \mathbb{U}_A \times \Omega$, whereas, in §9.2.2, it is described by an information σ-field $\mathcal{I}_\alpha \subset \mathcal{H}$. When this σ-field is a π-field (for instance, when $\mathbb{H}$ is finite), the connection between both approaches is given by

$$\mathcal{I}_\alpha = \pi(o_\alpha) \; , \tag{10.4}$$

where the π-field generated by a mapping has been introduced in Definition 3.32. $\lozenge$

Definition 10.3. *A* multi-agent stochastic input-output system *(MASIOS) is a collection consisting of agents A, states of Nature Ω and complete π-field $\mathcal{F}$, control sets and complete π-fields $\{\mathbb{U}_\alpha, \mathcal{U}_\alpha\}_{\alpha \in A}$, and observation functions $\{\mathbb{U}_\alpha, o_\alpha\}_{\alpha \in A}$.*

Example 10.4. For instance, if, in the description of a sequential optimal stochastic control problem as revealed in §4.5.1, informations fields $\mathcal{I}_t$ are given by signals $\boldsymbol{Y}_t : \mathbb{H} \to \mathbb{Y}_t$, where $(\mathbb{Y}_t, \mathcal{Y}_t)$ is some measurable space (see Remark 9.8 and Equation (10.4)), we obtain a MASIOS with agents $A = \{0, \dots, T-1\}$. △

Example 10.5. As another example of MASIOS, consider a state model as defined in §4.4.1. We set

$$A = \{0, \dots, T\} \; , \quad \Omega = \mathbb{X}_0 \times \prod_{t=1}^{T} \mathbb{W}_t \; , \tag{10.5}$$

so that states of Nature are scenarios

$$\omega = \big(x_0, w(\cdot)\big) = \big(x_0, w_1, w_2, \dots, w_T\big) \; . \tag{10.6}$$

Identifying any $u \in \mathbb{U}_{\{0,\dots,T\}}$ with an open-loop feedback $\gamma \equiv u$, we now define different observation functions as follows, with the help of the state map X_f (see Definition 4.7).

When the state x_t is observed at time t, this corresponds to

$$o_t(u, \omega) = X_f[0, x_0, u, w(\cdot)]_t \; . \tag{10.7}$$

The case when past states $x_0, \dots, x_t$ are observed at time t is given by

$$o_t(u, \omega) = \big(X_f[0, x_0, u, w(\cdot)]_0, \dots, X_f[0, x_0, u, w(\cdot)]_t\big) \; . \tag{10.8}$$

△

10.2.2 Control Laws

We define the counterpart of a policy in the Witsenhausen intrinsic model in §9.2: it is a control law, that is, a random variable defined on the universe Ω. An admissible control law for a focal agent is one that is measurable w.r.t. the agent closed-loop observation after control (of all agents).

Definition 10.6. *A* control law *for agent α is a random variable $\boldsymbol{U}_\alpha : \Omega \to \mathbb{U}_\alpha$, and a* collection of control laws *is a collection $\{\boldsymbol{U}_\beta\}_{\beta \in A}$ where $\boldsymbol{U}_\beta : \Omega \to \mathbb{U}_\beta$. We define the* set of collections of control laws *by:*

$$\mathcal{U}_A := \prod_{\beta \in A} \mathbb{U}_\beta^\Omega = \Big\{ \boldsymbol{U} = \{\boldsymbol{U}_\beta\}_{\beta \in A} \;\Big|\; \boldsymbol{U}_\beta : \Omega \to \mathbb{U}_\beta \; , \;\; \forall \beta \in A \Big\} \; . \tag{10.9}$$

We warn the reader that, though typographically close, the notations $\mathcal{U}_A$ for the set of collections of control laws in (10.9) and $\mathfrak{U}_A$ for the decision field in (10.1) are distinct. In what follows, except in Remark 10.10, we do not use the decision field notation.

Remark 10.7. Both $\mathbb{U}_\alpha$ and Ω being equipped with complete π-fields, control laws are necessarily measurable. Notice that a control law for agent α is a mapping defined over the universe Ω, whereas in §9.2.2 a policy for agent α was represented by a mapping $\lambda_\alpha : \mathbb{H} \to \mathbb{U}_\alpha$. A parallel can be established between a collection $\{\boldsymbol{U}_\beta\}_{\beta\in A}$ of control laws and the mapping $M_\lambda : \Omega \to \mathbb{U}_A$ attached to a collection $\lambda \in \Lambda_A^{\mathrm{ad}}$ of admissible policies, when the solvability property holds true (see Definition 9.10). In this chapter, control laws are random variables (see Definition 3.44). ♢

Definition 10.8. *For any collection $\boldsymbol{U} \in \mathcal{U}_A$ of control laws and for any agent $\alpha \in A$, the* observation of agent α after control *is the random variable $\eta_\alpha^{\boldsymbol{U}} : \Omega \to \mathbb{Y}_\alpha$ defined by*

$$\eta_\alpha^{\boldsymbol{U}}(\omega) := o_\alpha\big(\boldsymbol{U}(\omega), \omega\big) \ , \ \ \forall \omega \in \Omega \ . \tag{10.10}$$

The collection $\{\eta_\beta^{\boldsymbol{U}}\}_{\beta\in A}$ of random variables is called closed-loop observations.

In general, the observation available to agent α depends, through the collection $\boldsymbol{U} = \{\boldsymbol{U}_\beta\}_{\beta\in A}$ of control laws, upon the control laws of other agents by expanding (10.10) into

$$\eta_\alpha^{\boldsymbol{U}}(\omega) = o_\alpha\big(\{\boldsymbol{U}_\beta(\omega)\}_{\beta\in A}, \omega\big) \ . \tag{10.11}$$

A control law is said to be *admissible* for an agent if she makes her decision with no more than her observation after control.

Definition 10.9. *An* admissible control law *for agent α is a control law $\boldsymbol{U}_\alpha : \Omega \to \mathbb{U}_\alpha$ such that*

$$\boldsymbol{U}_\alpha \preceq \eta_\alpha^{\boldsymbol{U}} \ . \tag{10.12}$$

The set of admissible (collections of) control laws *is defined by:*

$$\mathcal{U}_A^{\mathrm{ad}} := \Big\{ \boldsymbol{U} = \{\boldsymbol{U}_\alpha\}_{\alpha\in A} \in \mathcal{U}_A \ \Big| \ \boldsymbol{U}_\alpha \preceq \eta_\alpha^{\boldsymbol{U}} \ , \ \ \forall \alpha \in A \Big\} \ . \tag{10.13}$$

The measurability constraint $\boldsymbol{U}_\alpha \preceq \eta_\alpha^{\boldsymbol{U}}$ is taken in the sense of measurability with respect to partition fields as in Definition 3.32 (see also Proposition 3.35).

Remark 10.10. Here, admissible control laws and the measurability constraint $\boldsymbol{U}_\alpha \preceq \eta_\alpha^{\boldsymbol{U}}$ are the counterparts of admissible policies $\lambda_\alpha : \mathbb{H} \to \mathbb{U}_\alpha$ measurable w.r.t. $\mathfrak{I}_\alpha$, that is, satisfying $\lambda_\alpha^{-1}(\mathfrak{U}_\alpha) \subset \mathfrak{I}_\alpha$ as in Definition 9.6. ♢

Remark 10.11. When not specified, the notation $\preceq$ is relative to mappings with common domain Ω (see §3.4.2). $\Diamond$

A special class of admissible control laws is the one made of *open-loop* or *deterministic* or *constant* control laws.

Definition 10.12. *The set* $\perp_A$ *of* open-loop control laws*, or* deterministic control laws*, consists of the constant control laws, namely control laws measurable w.r.t. the trivial* π*-field* $\{\emptyset, \Omega\}$ *on* Ω*:*

$$\perp_A := \left\{ \boldsymbol{U} = \{\boldsymbol{U}_\alpha\}_{\alpha \in A} \in \mathcal{U}_A \;\middle|\; \boldsymbol{U}_\alpha \preceq \{\emptyset, \Omega\} \,, \;\; \forall \alpha \in A \right\} \,. \tag{10.14}$$

Each $\boldsymbol{U}_\alpha : \Omega \to \mathbb{U}_\alpha$ in $\boldsymbol{U} = \{\boldsymbol{U}_\alpha\}_{\alpha \in A} \in \perp_A$ takes a constant value in $\mathbb{U}_\alpha$. The notation $\perp_A$ refers to the fact that the class of constant mappings is the bottom of the lattice of equivalence classes of mappings (see Proposition 3.42).

10.2.3 Precedence and Memory-Communication Relations

Thanks to the connection (10.4) between information fields and observations, we can characterize, in the MASIOS framework of §10.2, the precedence and memory-communication binary relations already introduced in §9.4.

For this purpose, we make use of the following notations. Consider $B \subset A$ a subset of agents. We set

$$u_B := \{u_\beta\}_{\beta \in B} \,, \tag{10.15}$$

and, for any collection $\{H_\alpha\}_{\alpha \in A}$ of mappings defined over Ω,

$$H_B := \{H_\beta\}_{\beta \in B} \,. \tag{10.16}$$

The *precedence* binary relation $\mathfrak{P}$ of Definition 9.15 identifies couples of agents, where the decision of the first agent indeed influences the observation of the second. By the correspondance (10.4), the subset $\langle\alpha\rangle_{\mathfrak{P}}$ of predecessors of α is (the smallest subset) such that there exist a mapping $\widetilde{o}_\alpha$ satisfying

$$o_\alpha(u, \omega) = \widetilde{o}_\alpha(u_{\langle\alpha\rangle_{\mathfrak{P}}}, \omega) \,, \tag{10.17}$$

expressing that $o_\alpha(u, \omega)$ depends only on the components $u_{\langle\alpha\rangle_{\mathfrak{P}}} = \{u_\beta\}_{\beta \in \langle\alpha\rangle_{\mathfrak{P}}}$ of the decision u.

The *memory-communication* binary relation $\mathfrak{M}$ of Definition 9.32 identifies couples of agents, where the observation of the first one is passed on to the second one. By the correspondance (10.4), the subset $\langle\alpha\rangle_{\mathfrak{M}}$ of agents whose information is embedded within the information of agent α is (the largest subset) such that:

$$o_{\langle\alpha\rangle_{\mathfrak{M}}}(\cdot,\cdot) \preceq_{\mathbb{U}_A \times \Omega} o_\alpha(\cdot,\cdot) \,, \;\; \forall \alpha \in A \,. \tag{10.18}$$

Here, we specify that measurability is w.r.t. to mappings with domain $\mathbb{U}_A \times \Omega$ (see Remark 10.11).

10.2.4 A Typology of MASIOS

Thanks to the precedence and information-memory relations, we now introduce a typology of MASIOS, inspired from the discussion in §9.5.1.

Partially Nested MASIOS

We say that a MASIOS is *partially nested* when the precedence relation $\mathfrak{P}$ is included in the memory-communication relation $\mathfrak{M}$, that is, when

$$\langle \alpha \rangle_{\mathfrak{P}} \subset \langle \alpha \rangle_{\mathfrak{M}} \ , \ \ \forall \alpha \in A \ , \tag{10.19}$$

or, by (10.18), when

$$o_{\langle \alpha \rangle_{\mathfrak{P}}}(\cdot, \cdot) \preceq_{\mathbb{U}_A \times \Omega} o_\alpha(\cdot, \cdot) \ , \ \ \forall \alpha \in A \ . \tag{10.20}$$

Remark 10.13. A consequence of Proposition 3.39 and of (10.20) is that, for all $\boldsymbol{U} \in \mathcal{U}_A$, we have that $o_{\langle \alpha \rangle_{\mathfrak{P}}}\big(\boldsymbol{U}(\cdot), \cdot\big) \preceq o_\alpha\big(\boldsymbol{U}(\cdot), \cdot\big)$ for all $\boldsymbol{U} \in \mathcal{U}_A$. This property is taken as the definition of a *partially nested information structure* in [80, 82]: it imposes conditions on the closed-loop observations (10.10), so that measurability is w.r.t. to mappings with domain Ω (see Remark 10.11). On the contrary, assumption (10.20) is an "open-loop" assumption, which does not require assumptions w.r.t. the closed-loop observations, and which makes use of measurability w.r.t. to mappings with domain $\mathbb{U}_A \times \Omega$. $\Diamond$

Sequential MASIOS

Consider the case where each agent in A is supposed to represent a time period t:

$$A = \{0, \dots, T\} \ \text{ where } \ T \in \mathbb{N}^* \ . \tag{10.21}$$

With the notations of §10.2.3, and especially Equation (10.1), we have that

$$\prod_{t=0}^{T} \mathbb{U}_t = \mathbb{U}_{\{0,\dots,T\}} \ . \tag{10.22}$$

Following §9.5.1, the MASIOS given by the family $\{o_t\}_{t=0,\dots,T}$ of observation functions

$$o_t : \mathbb{U}_{\{0,\dots,T\}} \times \Omega \to \mathbb{Y}_t \tag{10.23}$$

is said to be a *sequential MASIOS* if it is sequential with the ordering $0, \dots, T$. By (9.50) and with the notations of §10.2.3, this is equivalent to

$$\langle 0 \rangle_{\mathfrak{P}} = \emptyset \ \text{ and } \ \langle t \rangle_{\mathfrak{P}} \subset \{0, \dots, t-1\} \ , \ \ \forall t \in \{1, \dots, T\} \ . \tag{10.24}$$

In other words, the observation at time t depends at most upon the past decisions $u_0, \dots, u_{t-1}$ (and the state of Nature ω). Indeed, by (10.17) and Proposition 3.38, there exist mappings $\widetilde{o}_t$, for $t = 0, \dots, T$, such that

$$o_t(u_0,\dots,u_T,\omega)=\widetilde{o}_t(u_0,\dots,u_{t-1},\omega)\ ,\ \ \forall t=1,\dots,T\ , \tag{10.25}$$

with the special case $o_0(u_0,\dots,u_T,\omega)=\widetilde{o}_0(\omega)$.

Example 10.14. Following Example 10.5, consider the MASIOS induced by a state model as defined in §4.4.1. It can be checked that any expression of the form (with X_f the state map of Definition 4.7)

$$o_t(u,\omega)=\widetilde{o}_t\big(X_f[0,x_0,u,w(\cdot)]_0,\dots,X_f[0,x_0,u,w(\cdot)]_t,w(\cdot)\big) \tag{10.26}$$

defines a sequential MASIOS. This includes imperfect and corrupted observations of the past states.

An important class of sequential MASIOS is given by linear state models with linear observations. More precisely, linear state models are those for which the dynamics $f_t:\mathbb{X}_t\times\mathbb{U}_t\times\mathbb{W}_{t+1}\to\mathbb{X}_{t+1}$ are linear mappings. Linear observations correspond to $o_t(u,\omega)$ being a linear expression in $w(\cdot),X_f[0,x_0,u,w(\cdot)]_0,\dots,X_f[0,x_0,u,w(\cdot)]_t$. △

Quasiclassical MASIOS

As in §9.5.1, we say that a MASIOS is *quasiclassical* if it is sequential (with the ordering $0,\dots,T$) and partially nested, that is,

$$\langle 0\rangle_{\mathfrak{P}}=\emptyset\ \text{ and }\ \langle t\rangle_{\mathfrak{P}}\subset\{0,\dots,t-1\}\cap\langle t\rangle_{\mathfrak{M}}\ ,\ \ \forall t\in\{1,\dots,T\}\ . \tag{10.27}$$

In other words, if decisions made at time s affect the observation o_t ($s\in\langle t\rangle_{\mathfrak{P}}$), then $s\le t-1$ and the observation o_s is embedded in the observation o_t. Indeed, by $s\in\langle t\rangle_{\mathfrak{M}}$, (10.18) and Proposition 3.38, there exists a mapping $\mathfrak{f}_{s,t}$ such that $o_s=\mathfrak{f}_{s,t}(o_t)$.

Classical MASIOS

We say that a sequential MASIOS (with the ordering $0,\dots,T$) displays *perfect memory* if

$$\{0,\dots,t\}\subset\langle t\rangle_{\mathfrak{M}}\ ,\ \ \forall t\in\{0,\dots,T\}\ . \tag{10.28}$$

Information accumulates with time, that is, the π-fields $\mathfrak{I}_t=\pi(o_t)$ form a filtration:

$$\mathfrak{I}_0\subset\cdots\subset\mathfrak{I}_{t-1}\subset\mathfrak{I}_t\subset\cdots\subset\mathfrak{I}_T\ . \tag{10.29}$$

As a consequence of (9.52), a sequential MASIOS displaying perfect memory is classical (see §9.5.1), hence quasiclassical, with the ordering $0,\dots,T$.

Remark 10.15. Define *closed-loop perfect memory* as the property that, for all $t=0,\dots,T$,

$$o_t\big(\boldsymbol{U}(\cdot),\cdot\big)\preceq_{\Omega}o_{t+1}\big(\boldsymbol{U}(\cdot),\cdot\big)\ ,\ \ \forall\boldsymbol{U}\in\mathcal{U}_A^{\mathrm{ad}}\ . \tag{10.30}$$

As in Remark 10.13, this definition imposes conditions on the closed-loop observations (10.10), so that measurability is w.r.t. to mappings with domain Ω (see Remark 10.11). In contrast, Assumption (10.28) for the definition of perfect memory is an "open-loop" assumption, which does not require assumptions w.r.t. the closed-loop observations, and which makes use of measurability w.r.t. to mappings with domain $\mathbb{U}_{\{0,\dots,T\}} \times \Omega$. Open-loop perfect memory (10.28) implies closed-loop perfect memory (10.30). Notice that a weaker form of open-loop perfect memory, namely

$$o_t(u,\cdot) \preceq_\Omega o_{t+1}(u,\cdot) \ , \quad \forall u \in \mathbb{U}_{\{0,\dots,T\}} \ ,$$

does not imply closed-loop perfect memory (10.30). This can directly be seen with the following example: let $o_0(\omega) = \omega$ and $o_1(u,\omega) = u - \omega$; then $o_0(\cdot) \preceq o_1(u,\cdot)$ for all u; whereas, for $\boldsymbol{U}(\omega) = \omega$, this $\boldsymbol{U}$ is admissible since $\boldsymbol{U} \preceq o_0$, but obviously $o_0(\cdot) \not\preceq o_1\big(\boldsymbol{U}(\cdot),\cdot\big)$ since the latter is the zero mapping. $\Diamond$

10.3 No Open-Loop Dual Effect and No Dual Effect Control Laws

A collection of control laws induces a partition of the universe Ω. We say that no open-loop dual effect holds true when all constant control laws induce the same fixed partition.

10.3.1 No Open-Loop Dual Effect (NOLDE)

We now introduce the notion of *no open-loop dual effect*. For this purpose, we use the measurability equivalence $\equiv$ between mappings of Definition 3.40

Definition 10.16. *The property of* no open-loop dual effect *(NOLDE) holds true for the MASIOS discussed in §10.2 if we have that:*

$$\eta_\alpha^{\boldsymbol{U}} \equiv \eta_\alpha^{\boldsymbol{U}'} \ , \quad \forall (\boldsymbol{U},\boldsymbol{U}') \in \perp_A \times \perp_A \ , \quad \forall \alpha \in A \ . \tag{10.31}$$

In the case of NOLDE, for any agent $\alpha \in A$, all observations after open-loop control are equivalent, in the sense of measurability equivalence between mappings of Definition 3.40. Therefore, all observations after open-loop control are equivalent to a fixed mapping[1] ζ_α with domain Ω:

$$\eta_\alpha^{\boldsymbol{U}} \equiv \zeta_\alpha \ , \quad \forall \boldsymbol{U} \in \perp_A \ . \tag{10.32}$$

Example 10.17. It is shown in [127] that linear state models with linear observations, as defined in Example10.14, possess the NOLDE property if they display perfect memory as defined in §10.2.4. $\triangle$

[1] For instance, take for ζ_α any mapping of the class of $\eta_\alpha^{\boldsymbol{U}}$ for $\boldsymbol{U} \in \perp_A$.

Remark 10.18. In [81], Equation (5) expresses a similar property. ◊

The following proposition, adapted from [13], is a straightforward consequence of Proposition 3.41.

Proposition 10.19. *The property of no open-loop dual effect (NOLDE) holds true if, and only if, there exist a collection of mappings* $\{\mathfrak{f}_\alpha\}_{\alpha \in A}$ *where* $\mathfrak{f}_\alpha : \mathbb{U}_A \times \mathbb{Z}_\alpha \to \mathbb{Y}_\alpha$ *and a collection* $\{\zeta_\alpha\}_{\alpha \in A}$ *of random variables where* $\zeta_\alpha : \Omega \to \mathbb{Z}_\alpha$ *such that*

- *the partial mapping* $\mathfrak{f}_\alpha(u, \cdot) : \mathbb{Z}_\alpha \to \mathbb{Y}_\alpha$ *is injective, for all* $u \in \mathbb{U}_A$*;*
- *the observations satisfy* $o_\alpha(u, \omega) = \mathfrak{f}_\alpha\big(u, \zeta_\alpha(\omega)\big)$*, for all* $(u, \omega) \in \mathbb{U}_A \times \Omega$*.*

10.3.2 No Dual Effect Control Laws

No dual effect control laws are those control laws for which, in case of NOLDE, the closed-loop observations induce the same partitions as the constant control laws.

Definition 10.20. *Assume that the NOLDE property holds true, with the fixed observations* ζ *as in* (10.32)*. The* no dual effect control laws set *is made of all admissible control laws such that the closed-loop observations are equivalent to the fixed mapping* ζ_α*:*

$$\mathcal{U}_A^{\mathrm{nde}} := \Big\{ \boldsymbol{U} = \{\boldsymbol{U}_\alpha\}_{\alpha \in A} \in \mathcal{U}_A \;\Big|\; \eta_\alpha^{\boldsymbol{U}} \equiv \zeta_\alpha \;, \;\; \forall \alpha \in A \Big\} \cap \mathcal{U}_A^{\mathrm{ad}} \,. \tag{10.33}$$

Thus, “closing” the system with any control law belonging to the no dual effect control law set produces the same fixed closed-loop observations.

Definition 10.21. *Assume that the NOLDE property holds true, with the fixed observations* ζ *as in* (10.32)*. The* set of control laws measurable w.r.t. the fixed observations $\zeta = \{\zeta_\alpha\}_{\alpha \in A}$ *is defined by:*

$$\mathcal{U}_A^{\zeta} := \big\{ \boldsymbol{U} = \{\boldsymbol{U}_\alpha\}_{\alpha \in A} \in \mathcal{U}_A \;\big|\; \boldsymbol{U}_\alpha \preceq \zeta_\alpha \;, \;\; \forall \alpha \in A \big\} \;. \tag{10.34}$$

We have the following relation between the no dual effect control laws set $\mathcal{U}_A^{\mathrm{nde}}$ in (10.33) and the set $\mathcal{U}_A^{\zeta}$ in (10.34).

Proposition 10.22. *Assume that the NOLDE property holds true, with the fixed observations* ζ *as in* (10.32)*. Then, no dual effect control laws are necessarily measurable w.r.t. the fixed observation* ζ*, that is,*

$$\mathcal{U}_A^{\mathrm{nde}} \subset \mathcal{U}_A^{\zeta} \,. \tag{10.35}$$

Proof. Let $\boldsymbol{U} = \{\boldsymbol{U}_\alpha\}_{\alpha \in A} \in \mathcal{U}_A^{\mathrm{nde}}$. On the one hand, we have that $\boldsymbol{U}_\alpha \preceq \eta_\alpha^{\boldsymbol{U}}$, for all agent $\alpha \in A$, since $\boldsymbol{U} \in \mathcal{U}_A^{\mathrm{ad}}$ by (10.33) and (10.13). On the other hand, we have that $\eta_\alpha^{\boldsymbol{U}} \equiv \zeta_\alpha$ by (10.33) and (10.32). Thus, $\boldsymbol{U}_\alpha \preceq \eta_\alpha^{\boldsymbol{U}} \equiv \zeta_\alpha$. Since this holds true for any agent α, we conclude that $\boldsymbol{U} \in \mathcal{U}_A^{\zeta}$. □

10.3.3 Characterization of No Dual Effect Control Laws

We now characterize the no dual effect control laws according to the typology discussed in §10.2.4. We use the following two lemmas.

Lemma 10.23. *Consider three mappings* $H_i \colon \Omega \to \mathbb{Y}_i$, $i = 1, 2$ *and* $f \colon \mathbb{Y}_1 \times \Omega \to \mathbb{Y}_3$. *Assume that, for all* $y_1 \in \mathbb{Y}_1$, $f(y_1, \cdot) \preceq H_2(\cdot)$ *and that* $H_1(\cdot) \preceq H_2(\cdot)$. *Then* $f\big(H_1(\cdot), \cdot\big) \preceq H_2(\cdot)$.

Proof. Let $(\omega, \omega') \in \Omega^2$ be such that $H_2(\omega) = H_2(\omega')$. Since $H_1(\cdot) \preceq H_2(\cdot)$, we have that $H_1(\omega) = H_1(\omega')$ by Proposition 3.38. Putting $y_1 = H_1(\omega) = H_1(\omega')$, we thus get $f(y_1, \omega) = f(y_1, \omega')$ since $f(y_1, \cdot) \preceq H_2(\cdot)$. We conclude that

$$f\big(H_1(\omega), \omega\big) = f(y_1, \omega) = f(y_1, \omega') = f\big(H_1(\omega'), \omega'\big) \; .$$

The proof is complete by Proposition 3.38. □

Lemma 10.24. *Let* $H_i \colon \Omega \to \mathbb{Y}_i$, $i = 1, 2$ *and* $f \colon \mathbb{Y}_1 \times \Omega \to \mathbb{Y}_3$. *Assume that, for all* $y_1 \in \mathbb{Y}_1$, $H_2(\cdot) \preceq f(y_1, \cdot)$, *and that* $H_1(\cdot) \preceq f\big(H_1(\cdot), \cdot\big)$. *Then* $H_2(\cdot) \preceq f\big(H_1(\cdot), \cdot\big)$.

Proof. Let $(\omega, \omega') \in \Omega^2$ be such that $f\big(H_1(\omega), \omega\big) = f\big(H_1(\omega'), \omega'\big)$. Since $H_1(\cdot) \preceq f\big(H_1(\cdot), \cdot\big)$, we have that $H_1(\omega) = H_1(\omega')$. Putting $y_1 = H_1(\omega) = H_1(\omega')$, we thus get

$$f(y_1, \omega) = f\big(H_1(\omega), \omega\big) = f\big(H_1(\omega'), \omega'\big) = f(y_1, \omega') \; .$$

On the other hand, we have that $H_2(\cdot) \preceq f(y_1, \cdot)$, so that $H_2(\omega) = H_2(\omega')$. The proof is complete by Proposition 3.38. □

Partially Nested MASIOS

The following main result, established in [13], provides a description of the set of no dual effect control laws for MASIOS displaying the NOLDE property.

Theorem 10.25 ([13]). *Assume that the NOLDE property holds true, with the fixed observations* ζ *as in* (10.32). *Assume that the MASIOS is partially nested as in* (10.20). *Then, the no dual effect control laws in* (10.33) *are exactly the admissible control laws which are measurable w.r.t. the fixed observation* ζ:

$$\mathcal{U}_A^{\mathrm{nde}} = \mathcal{U}_A^{\mathrm{ad}} \cap \mathcal{U}_A^{\zeta} \; . \tag{10.36}$$

Proof. By Proposition 10.22, it suffices to show that $\mathcal{U}_A^{\mathrm{ad}} \cap \mathcal{U}_A^{\zeta} \subset \mathcal{U}_A^{\mathrm{nde}}$.

Let $\boldsymbol{U} = \{\boldsymbol{U}_\beta\}_{\beta \in A} \in \mathcal{U}_A^{\mathrm{ad}} \cap \mathcal{U}_A^{\zeta}$, that is,

$$\boldsymbol{U}_\beta \preceq \zeta_\beta \text{ and } \boldsymbol{U}_\beta \preceq \eta_\beta^{\boldsymbol{U}} \; , \; \forall \beta \in A \; . \tag{10.37}$$

Let $\alpha \in A$ be fixed: we now prove that both $\eta_\alpha^{\boldsymbol{U}} \preceq \zeta_\alpha$ and $\zeta_\alpha \preceq \eta_\alpha^{\boldsymbol{U}}$ hold true. We use the property that, by Equation (10.17) and by abuse of notation (see (10.16)),

$$\eta_\alpha^{\boldsymbol{U}} = \eta_\alpha^{\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}}} \ , \ \forall \alpha \in A \ . \tag{10.38}$$

First, we show that $\eta_\alpha^{\boldsymbol{U}} \preceq \zeta_\alpha$. For any $u \in \mathbb{U}_A$ (identified with a constant control law), we have that:

$$\begin{aligned}
\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}} &= \{\boldsymbol{U}_\beta\}_{\beta\in\langle\alpha\rangle_{\mathfrak{P}}} && \text{by (10.16)}\\
&\equiv \bigvee_{\beta\in\langle\alpha\rangle_{\mathfrak{P}}} \boldsymbol{U}_\beta && \text{by Proposition 3.42}\\
&\preceq \bigvee_{\beta\in\langle\alpha\rangle_{\mathfrak{P}}} \zeta_\beta && \text{because } \boldsymbol{U}_\beta \preceq \zeta_\beta \text{ by (10.37)}\\
&\equiv \bigvee_{\beta\in\langle\alpha\rangle_{\mathfrak{P}}} o_\beta(u,\cdot) \ , \ \forall u \in \mathbb{U}_A && \text{by (10.31), (10.32) and (10.10)}\\
&\equiv \{o_\beta(u,\cdot)\}_{\beta\in\langle\alpha\rangle_{\mathfrak{P}}} && \text{by Proposition 3.42}\\
&\equiv o_{\langle\alpha\rangle_{\mathfrak{P}}}(u,\cdot) && \text{by (10.16)}\\
&\preceq o_\alpha(u,\cdot) && \text{by (10.20)}\\
&\equiv \zeta_\alpha(\cdot) && \text{by (10.32).}
\end{aligned}$$

Therefore, we have that, on the one hand, $\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}}(\cdot) \preceq o_\alpha(u,\cdot) \equiv \zeta_\alpha(\cdot)$ and, on the other hand, $o_\alpha(u,\cdot) = \widetilde{o}_\alpha(u_{\langle\alpha\rangle_{\mathfrak{P}}},\cdot) \equiv \zeta_\alpha(\cdot)$, for all $u \in \mathbb{U}_A$ by (10.17) and (10.32). By Lemma 10.23, we deduce that:

$$\eta_\alpha^{\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}}}(\cdot) = \widetilde{o}_\alpha\big(\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}}(\cdot),\cdot\big) \preceq \zeta_\alpha(\cdot) \ .$$

By (10.38), we conclude that

$$\eta_\alpha^{\boldsymbol{U}} \preceq \zeta_\alpha \ .$$

Second, we prove that $\zeta_\alpha \preceq \eta_\alpha^{\boldsymbol{U}}$. We have that

$$
\begin{aligned}
\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}} &= \{\boldsymbol{U}_\beta\}_{\beta\in\langle\alpha\rangle_{\mathfrak{P}}} && \text{by (10.16)}\\
&\equiv \bigvee_{\beta\in\langle\alpha\rangle_{\mathfrak{P}}} \boldsymbol{U}_\beta && \text{by Proposition 3.42}\\
&\preceq \bigvee_{\beta\in\langle\alpha\rangle_{\mathfrak{P}}} \eta_\beta^{\boldsymbol{U}} && \text{because } \boldsymbol{U}_\beta \preceq \eta_\beta^{\boldsymbol{U}} \text{ by (10.37)}\\
&\equiv \eta_{\langle\alpha\rangle_{\mathfrak{P}}}^{\boldsymbol{U}} && \text{by (10.16) and Proposition 3.42}\\
&= o_{\langle\alpha\rangle_{\mathfrak{P}}}\big(\boldsymbol{U}(\cdot),\cdot\big) && \text{by (10.10)}\\
&\preceq o_\alpha\big(\boldsymbol{U}(\cdot),\cdot\big) && \text{by the partially nested property (10.20)}\\
&= \widetilde{o}_\alpha\big(\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}}(\cdot),\cdot\big) && \text{by (10.17).}
\end{aligned}
$$

Therefore, on the one hand, we have just proven that

$$\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}}(\cdot) \preceq \widetilde{o}_\alpha\big(\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}}(\cdot),\cdot\big) \ .$$

On the other hand, for all $u \in \mathbb{U}_A$, we have that

$$
\begin{aligned}
\zeta_\alpha(\cdot) &\equiv \eta_\alpha^u(\cdot) && \text{by (10.32)}\\
&= o_\alpha(u,\cdot) && \text{by (10.10)}\\
&= \widetilde{o}_\alpha(u_{\langle\alpha\rangle_{\mathfrak{P}}},\cdot) && \text{by (10.17).}
\end{aligned}
$$

By Lemma 10.24, we deduce that

$$\zeta_\alpha(\cdot) \preceq \widetilde{o}_\alpha\big(\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}}(\cdot),\cdot\big) \ .$$

By (10.38), we conclude that

$$\zeta_\alpha \preceq \eta_\alpha^{\boldsymbol{U}_{\langle\alpha\rangle_{\mathfrak{P}}}} = \eta_\alpha^{\boldsymbol{U}} \ .$$

This completes the proof. □

Quasiclassical MASIOS

For sequential MASIOS, where each agent is supposed to represent a time period, we are able to obtain a result, that is more precise than Theorem 10.25. Indeed, we now show that, for quasiclassical MASIOS displaying the NOLDE property, the no dual effect control laws are the control laws which are measurable w.r.t. the fixed observations.

Proposition 10.26. *Assume that the NOLDE property holds true, with the fixed observations ζ as in (10.32). Assume that the MASIOS is quasiclassical, as in 10.2.4 with the ordering $0,\dots,T$ of agents. Then, the no dual effect* *control laws in* (10.33) *are the control laws which are measurable w.r.t. the fixed observation* ζ*, that is,*

$$\mathcal{U}^{\mathrm{nde}}_{\{0,\dots,T\}} = \mathcal{U}^{\zeta}_{\{0,\dots,T\}} \ . \tag{10.39}$$

Proof. By Proposition 10.22, it suffices to show that $\mathcal{U}^{\zeta}_{\{0,\dots,T\}} \subset \mathcal{U}^{\mathrm{nde}}_{\{0,\dots,T\}}$. Let $\boldsymbol{U} = \{\boldsymbol{U}_t\}_{t=0,\dots,T} \in \mathcal{U}^{\zeta}_{\{0,\dots,T\}}$, that is,

$$\boldsymbol{U}_t \preceq \zeta_t \ , \ \ \forall t = 0,\dots,T \ . \tag{10.40}$$

We prove by induction that

$$\Big(\forall t = 0,\dots,T \ , \ \ \boldsymbol{U}_t \preceq \zeta_t\Big) \Rightarrow \Big(\forall t = 0,\dots,T \ , \ \ \boldsymbol{U}_t(\cdot) \preceq o_t\big(\boldsymbol{U}(\cdot),\cdot\big)\Big) \ .$$

Let the induction assumption $\mathrm{H}(t)$ be

$$\Big(\forall s = 0,\dots,t \ , \ \ \boldsymbol{U}_s \preceq \zeta_s\Big) \Rightarrow \Big(\forall s = 0,\dots,t \ , \ \ \boldsymbol{U}_s(\cdot) \preceq o_s\big(\boldsymbol{U}(\cdot),\cdot\big)\Big) \ .$$

Suppose that $\boldsymbol{U}_0(\cdot) \preceq \zeta_0(\cdot)$. By (10.31), we know that $\zeta_0(\cdot) \equiv o_0(u,\cdot)$, for all $u \in \mathbb{U}$. However, o_0 is independent of u, since agent 0 has no predecessor ($[0] = \emptyset$). Thus, we conclude that

$$\boldsymbol{U}_0(\cdot) \preceq o_0(u,\cdot) = o_0\big(\boldsymbol{U}(\cdot),\cdot\big)$$

and the induction assumption $\mathrm{H}(0)$ holds true.

Assume that the induction assumption $\mathrm{H}(t-1)$ holds true, and suppose that

$$\boldsymbol{U}_s \preceq \zeta_s \ , \ \ \forall s = 0,\dots,t \ . \tag{10.41}$$

We have that

$$\begin{aligned}
\boldsymbol{U}_{\langle t\rangle_{\mathfrak{P}}}(\cdot) &= \bigvee_{s\in\langle t\rangle_{\mathfrak{P}}} \boldsymbol{U}_s(\cdot) && \text{by (10.16)}\\
&\preceq \bigvee_{s\in\langle t\rangle_{\mathfrak{P}}} o_s\big(\boldsymbol{U}(\cdot),\cdot\big) && \text{by assumption } \mathrm{H}(t-1)\\
& && \text{and since } \langle t\rangle_{\mathfrak{P}} \subset \{0,\dots,t-1\} \text{ by (10.27)}\\
&\preceq \bigvee_{s\in\langle t\rangle_{\mathfrak{M}}} o_s\big(\boldsymbol{U}(\cdot),\cdot\big) && \text{by the partially nested property (10.19)}\\
&\equiv o_{\langle t\rangle_{\mathfrak{M}}}\big(\boldsymbol{U}(\cdot),\cdot\big) && \text{by (10.16)}\\
&\preceq o_t\big(\boldsymbol{U}(\cdot),\cdot\big) && \text{by (10.18)}\\
&= \widetilde{o}_t\big(\boldsymbol{U}_{\langle t\rangle_{\mathfrak{P}}}(\cdot),\cdot\big) && \text{by (10.17).}
\end{aligned}$$

Therefore, we have proved that

$$\boldsymbol{U}_{\langle t \rangle_{\mathfrak{P}}}(\cdot) \preceq \widetilde{o}_t\big(\boldsymbol{U}_{\langle t \rangle_{\mathfrak{P}}}(\cdot), \cdot\big) = o_t\big(\boldsymbol{U}(\cdot), \cdot\big) \ .$$

Now, on the other hand, we have that

$$\boldsymbol{U}_t(\cdot) \preceq \zeta_t(\cdot) \equiv o_t(u_{\langle t \rangle_{\mathfrak{P}}}, \cdot) \ , \ \ \forall \{u_s\}_{s \in \{0,\dots,T\}} \in \mathbb{U}_{\{0,\dots,T\}} \ ,$$

by (10.41), (10.31) and (10.17). We now use Lemma 10.24 with $H_1 = \boldsymbol{U}_t$ and $f(u, \omega) = o_t(u, \omega)$ to obtain that[2]

$$\boldsymbol{U}_t(\cdot) \preceq o_t\big(\boldsymbol{U}(\cdot), \cdot\big) \ . \tag{10.42}$$

Thus, assumption H(t) holds true. This completes the induction. □

Proposition 10.26 applies to sequential MASIOS displaying perfect memory because, as a consequence of (9.52), they are classical (see §9.5.1), hence quasiclassical, with the ordering $0, \dots, T$.

10.4 Conclusion

In this chapter, we have more deeply analyzed the "dual effect" of control previously discussed in §1.1.3, §1.2.1, §1.3.2 and §4.2.3. The specificity of sequential systems with perfect memory has been emphasized. When they display the NOLDE property, the no dual effect control laws have a simple characterization: they are the control laws which are measurable w.r.t. the fixed observations. Therefore, this chapter brings to light another element possibly explaining the importance of sequential systems with perfect memory in stochastic control.

[2] In fact, we use a slight variation of Lemma 10.24. Indeed $\boldsymbol{U}_{\langle t \rangle_{\mathfrak{P}}}(\cdot) \preceq o_t\big(\boldsymbol{U}_{\langle t \rangle_{\mathfrak{P}}}(\cdot), \cdot\big)$ is a weaker assumption than $\boldsymbol{U}(\cdot) \preceq o_t\big(\boldsymbol{U}(\cdot), \cdot\big)$ However, thanks to (10.17), the proof of Lemma 10.24 can easily be adapted to obtain the same conclusion.