

Informational Issues in Decentralized Control

Meir Pachter and Khanh Pham

Abstract The long-standing decentralized optimal control problem posed by Witsenhausen is analyzed and the underlying modeling issues are discussed. The strong connection to communication theory is highlighted. Informational aspects are emphasized and it is shown that, to some extent, Witsenhausen's decentralized optimal control problem is somewhat contrived.

Keywords Decentralized Control • Game Theory • Witsenhausen's Problem

1 Introduction

Informational issues in decentralized control are discussed. In this regard, Witsenhausen's problem [1] is the simplest decentralized optimal control problem. Indeed, the “simple” LQG control problem posed by Witsenhausen in his seminal paper made a great splash when it first appeared in 1968 because the optimal linear strategy is not optimal. This is caused by the nonclassical information pattern. Since then numerous attempts have been made in the intervening 45 years to obtain a better estimate of the minimal cost. A special session devoted to the Witsenhausen problem was held at the 2008 CDC and this stimulated renewed interest in the subject matter. Recently numerous additional papers on Witsenhausen's counterexample have appeared, and this problem was also extensively addressed in the most recent CDCs [2–6]. There is a fascination with Witsenhausen's counterexample in control circles and for good reason: It touches on the dual control issue where one needs to strike a balance between exploration and exploitation—this, exclusively due to the information pattern, which is non-nested. However, the objective of this article is not to survey the field, nor is it to further slightly improve the estimate of the minimal cost—the current best estimate of the minimal cost for

M. Pachter (✉)

Air Force Institute of Technology, Wright-Patterson A.F.B., OH 45433, USA

e-mail: meir.pachter@afit.edu

K. Pham

Air Force Research Laboratory, Kirtland A.F.B., NM 87117, USA

e-mail: khanh.pham@kirtland.af.mil

the “canonical” problem parameters currently stands at 0.1670790. The objective is to gain an understanding of the underlying engineering or physical problem that Witsenhausen’s mathematical model is addressing. In this respect, the coupled communications and control aspects of Witsenhausen’s problem are discussed and the attendant informational issues are carefully examined. Also, our aim is to gain physical insight into a range of methods for obtaining suboptimal solutions and, by doing so, dispel some of Witsenhausen’s counterexample mystique. The strong connection to communication theory is emphasized and the informational aspects of the problem are highlighted. The latter seem to direct one to the conclusion that Witsenhausen’s decentralized optimal control problem is to some extent contrived.

The paper is organized as follows. In Sect. 2 Witsenhausen’s problem is carefully stated. In Sect. 3 the communications aspect of Witsenhausen’s decentralized LQG optimal control problem are analyzed and the connection to detection theory is elucidated in Appendix A. The special case where the “receiver’s” noise floor is high is analyzed in Sect. 4 and the optimal modulation and detection scheme is shown to be linear. Concluding remarks are made in Sect. 5.

2 Witsenhausen’s Problem Statement

In Witsenhausen’s paper [1] the following decentralized LQG optimal control problem is considered.

1. *Dynamics*: The discrete-time dynamics are linear and scalar, and the planning horizon is $N = 2$. There are two “players,” Player 1 and Player 2. Player 1 acts at decision time $k = 0$ and his input is u_0 :

$$x_1 = x_0 + u_0, \quad x_0 \sim \mathcal{N}(0, \sigma_0^2). \quad (1)$$

Player 2 acts at decision time $k = 1$ and his input is u_1 :

$$x_2 = x_1 - u_1. \quad (2)$$

2. *Information*: The information of Player 1 at his decision time $k = 0$ is the initial state x_0 . The information available to Player 2 at his decision time $k = 1$ is a noise corrupted measurement of the state at time $k = 1$,

$$z_1 = x_1 + v_1, \quad v_1 \sim \mathcal{N}(0, \sigma^2). \quad (3)$$

Thus, the initial state x_0 is the private information of Player 1 which is not shared with Player 2. Player 1 has perfect information at his decision time $k = 0$, whereas Player 2 which acts at time $k = 1$ has access to the noise corrupted

measurement z_1 of the state x_1 . Thus, at his decision time $k = 1$, Player 2 has imperfect information on the state x_1 . In addition, Player 2 does not know the control u_0 of Player 1.

However, note that in [1], and in Eq. (1), it is also stated that the initial state x_0 is a random variable which is Gaussian distributed. This important point will be further discussed in the sequel.

3. *Strategy*: The so far specified information pattern mandates that the strategy of Player 1 is

$$u_0 = \gamma_0(x_0) \quad (4)$$

and the strategy of Player 2 is

$$u_1 = \gamma_1(z_1). \quad (5)$$

4. *Payoff*: The cost function, which both players strive to minimize, is

$$J(u_0, u_1; x_0) = K^2 u_0^2 + x_2^2. \quad (6)$$

Player 1, who has perfect state information, has a penalty K^2 on his control effort, whereas the control effort of Player 2, whose states' measurements are corrupted by noise, is free. Both players want the terminal state x_2 to be "small" while at the same time, the control effort, exclusively expended by Player 1, also to be "small." Evidently, Player 2 could effortlessly make the terminal state $x_2 \approx 0$, if only he knew the state x_1 with a high degree of precision.

As far as Player 1 is concerned, the random variable in the problem statement is the measurement error v_1 of Player 2. As far as Player 2 is concerned, the random variables in the problem statement are the initial state x_0 and his measurement error v_1 . The players are interested in the expectation of the cost function (6), which requires Player 1 to take the expectation on the random variable v_1 and Player 2 takes the expectation on the random variables v_1 and x_0 . Since the Players 1 and 2 have private information, x_0 and z_1 , respectively, their expected costs will be conditional on their private information. This brings us into the realm of nonzero-sum games.

Witsenhausen's decentralized control problem is schematically illustrated in Fig. 1. In Fig. 1 we have allowed for a more general initial state specification,

$$x_0 \sim \mathcal{N}(\bar{x}_0, \sigma_0^2) \quad (7)$$

and, without loss of generality, have set the parameter $\sigma = 1$; if $\sigma = 0$, the optimization problem is trivial—obviously, the optimal controls are $u_0^* = 0$ and $u_1^* = z_1$. The other extreme case where $\sigma \rightarrow \infty$ will be addressed in Sect. 4. Thus, the problem parameters are K and σ_0 .




Fig. 1 Witsenhausen's decentralized control problem

2.1 Initial State Information

In Witsenhausen's paper it is stated that the initial state x_0 is *random*, Gaussian distributed, and its statistics are specified according to Eq. (1), namely,

$$x_0 \sim \mathcal{N}(0, \sigma_0^2).$$

The informational aspect of the specification of the initial state's statistics merits special attention.

In Witsenhausen's paper [1], attention is confined to the special case where the statistic $\bar{x}_0 = 0$. It stands to reason that the consideration of a more general initial state is warranted, s.t. its statistics are specified by Eq. (7)—see also Fig. 1. Indeed, including in the problem statement the initial state's statistics immediately begs the question: why not allow for a more generic form of the initial state information, as specified in Eq. (7), whereupon the issue of who knows what, when, is immediately brought up; exclusively considering the special case where the initial state's statistic $\bar{x}_0 = 0$ tends to hide the importance of the \bar{x}_0 information. This brings to the forefront the question of whether Player 1, or Player 2, or both players are privy to the \bar{x}_0 information.

Consider Eqs. (7) or (1). This can be construed to mean that at time $k = 0$, Player 2 took a measurement of the initial state x_0 so that at his decision time $k = 1$, Player 2, in addition to the measurement z_1 , also has imperfect information about the state x_0 . Thus, the strategy (5) of Player 2 should in fact be replaced by the strategy

$$u_1 = \gamma_1(\bar{x}_0, z_1). \quad (8)$$

Whether or not Player 1 is informed on the initial state's measurement taken by Player 2 is an important question. If indeed the measurement information \bar{x}_0 is shared with Player 1 then his strategy (4) should be replaced by the strategy

$$u_0 = \gamma_0(x_0, \bar{x}_0). \quad (9)$$

Clearly, should at time $k = 0$ the information on the measurement \bar{x}_0 of Player 2 be shared with Player 1, Witsenhausen's control problem would be somewhat less decentralized than it appears to be at first glance.

Alternatively, one might argue that the specifications in Eqs. (1) or (7) mean that although it is so that at his decision time $k = 0$, Player 1 has perfect state information, he might be aware of the fact that the initial state x_0 presented to him at time $k = 0$ will actually be drawn from the distributions (1) or (7). This is critical information if a degree of information sharing among the players before the kickoff of this seemingly decentralized control problem is in fact allowed, or is required, to take place—in which case Player 1 employs a *prior commitment* strategy and the information \bar{x}_0 plays a crucial role in its synthesis. In this case, the strategy of Player 1 will take the somewhat unconventional form of Eq. (9). Concerning Eq. (9), one could then wonder why would Player 1 want to know the statistic of the initial state, since he already knows the initial state proper, but as already stated above, this additional information is required in order for Player 1 to be able to synthesize his prior commitment strategy. In this instance, whether or not also Player 2 is informed on the initial state's statistics will be discussed in the sequel. Suffice it to say that if Player 2 is informed about the statistic \bar{x}_0 , the strategy (5) of Player 2 will be replaced by the strategy (8). In this case the strategy (9) of Player 1, which acts at time $k = 0$, could be viewed as a delayed commitment strategy which encodes the fact that at his decision time $k = 0$, Player 1 knows that Player 2, who's turn will come at time $k = 1$, knows the statistic \bar{x}_0 of the initial state x_0 . In other words, both players know that the initial state x_0 will be drawn from a distribution (1) or (7), except that, prior to his move at time $k = 0$, Player 1 is given the initial state information.

When the initial state's measurement/information \bar{x}_0 is available to Player 2 and at time $k = 0$ this information is shared with Player 1, then Player 1, who has perfect information on the initial state x_0 , also knows that Player 2 knows that the initial state is distributed according to Eqs. (1) or (7). By virtue of this fact, the strategy of Player 1, which incorporates all the information available to him, takes the rather unusual form (9). This point, whether or not at his decision time $k = 0$ Player 1 knows that Player 2 knows that the initial state is distributed according to Eqs. (1) or (7), will be emphasized in Sect. 4.

Suppose the measurement (7) of the initial state taken by Player 2 is shared with Player 1, which now uses the strategy (9). In other words, the result of the initial state measurement of Player 2 is communicated to Player 1. Alternatively, suppose the information (1), or (7), of Player 1 concerning the p.d.f. wherefrom the initial state will be drawn is communicated to Player 2 before the start of the game. In both cases the information \bar{x}_0 is shared at time $k = 0$ and the respective strategies of Players 1 and 2 will be specified by Eqs. (9) and (8). This is the operational information pattern in Witsenhausen's paper. The control problem at hand now assumes a less decentralized character.

2.2 Cost Function

The ramifications as far as the cost functional is concerned are as follows. Since \bar{x}_0 is public information, the strategies' dependence on \bar{x}_0 is suppressed. The fact that the players have private information would naturally lead to a formulation where each player strives to minimize his own cost function: in the case of Player 1 it would be the expectation on v_1 of the cost function (6), conditional on his private information x_0 , and in the case of Player 2 it would be the expectation on x_0 and v_1 of the cost function (6), conditional on his private information z_1 . The players' strategies would be delayed commitment strategies, which means that the optimization of their respective cost functions would be performed in the Euclidean space R^1 .

Although the players have private information, it is nevertheless stipulated in [1] that both players are minimizing a common cost functional, namely the expectation on x_0 and v_1 of the cost function (6): according to Witsenhausen [1] and the many papers written on Witsenhausen's problem, the cost functional is

$$J(\gamma_0, \gamma_1; \bar{x}_0) = E_{x_0, v_1} (K^2(\gamma_0(x_0))^2 + (x_0 + \gamma_0(x_0) - \gamma_1(x_0 + \gamma_0(x_0) + v_1))^2). \quad (10)$$

This definition of the common cost functional is made possible by the fact that both players share the information on the statistics (1), or (7), of the initial state x_0 . Now, the players' strategies $\gamma_0(x_0)$ and $\gamma_1(z_1)$ are prior commitment strategies. Indeed, Player 2 now employs a prior commitment strategy—he decides on his optimal strategy $\gamma_1(\cdot)$ ahead of time, before ever receiving the measurement z_1 , which does not feature in Witsenhausen's cost functional. In order for the cost of Player 2 which employs a prior commitment strategy to equal the average realized cost of Player 1, the latter is minimizing the expectation of the cost (6), taken not just over the random variable v_1 but also over the initial state x_0 . This gives the appearance of Player 1 playing a *prior commitment* game, that is, instead of determining his control u_0 upon receiving his initial state information x_0 , Player 1 determines his *strategy* function $\gamma_0(\cdot)$ ahead of time. Indeed, his private information x_0 does not feature in Witsenhausen's cost functional. The “prior commitment” aspect of the strategy of Player 1 in this inherently cooperative game further manifests itself in the scenario discussed in Sect. 3 where *prior communication* is allowed, that is, the players are allowed to communicate before kickoff time, in which case Player 1 informs Player 2 about his optimal or suboptimal *strategy* prior to time $k = 0$. This requires Player 1 to be privy to the statistic \bar{x}_0 of the initial state x_0 , as indeed he must be if he is to minimize Witsenhausen's cost functional. Evidently, the requirement of prior communication further diminishes the much touted decentralized control character of Witsenhausen's problem.

The dynamic “game” is now in *normal form*, is static, and is not in *extensive form*. The price to pay for transforming the dynamic game into normal form is that prior commitment strategies are used. When the game is in extensive form and delayed

commitment strategies are used, part of the optimization is carried out in Euclidean space. Now that prior commitment strategies are used, the optimization must be performed in function space.

3 Communication and Control

To gain an understanding of, and insight into, the decentralized optimal control problem at hand, the communications context of Witsenhausen's decentralized control problem is now discussed. By directly viewing the decentralized optimal control problem (1)–(5), (7)–(10) as a communications problem, which indeed it is, and was originally perceived by Witsenhausen, one realizes that notwithstanding the fact that the cost function is quadratic in the controls, the problem at hand entails the minimization of a non-convex and complex cost functional, which is a hard problem. At the same time, a hierarchy of suboptimal solutions readily suggests itself. This must have been the motivation for Witsenhausen's original counterexample in the first place [1]. However, rather than directly viewing Witsenhausen's problem as an optimization problem in function space, we shall dwell on the physical meaning of the mathematical problem posed by Witsenhausen. By emphasizing the communications context of Witsenhausen's decentralized optimal control problem one opens wide the door to the synthesis of a family of suboptimal solutions of Witsenhausen's decentralized optimal control problem, as evidenced by the rich current literature. In doing so we hope to dispel some of Witsenhausen's counterexample mystique.

A communication problem is considered where both Player 1, the transmitter, and Player 2, the receiver, know that the "message" x_0 will be selected according to the probability distribution specified by Eq. (7). Player 1 will encode the information x_0 according to

$$x_1 = f(x_0)$$

before sending x_1 over a Gaussian communication channel whose statistics are specified by Eq. (3). The optimal function $f(\cdot)$ must be determined. It is also specified that the cost to Player 1 of encoding the information x_0 is $K^2(f(x_0) - x_0)^2$. Player 2, using his measurement z_1 , estimates the received signal x_1 . Player 2 strives to reduce the variance of the estimation error of x_1 . Both players want to minimize the average expected total cost of encoding and decoding the transmitted signal. As such, this formulation models a communication problem, but not a decentralized optimal control problem or a dynamic game.

1. In the context of a communication problem, we need to assume that the initial state's information/measurement (1), or (7) of Player 2, is shared with Player 1.

Thus, both players know that the initial state will be drawn from the Gaussian distribution (7), that is, the initial state's p.d.f. is

$$\phi(x_0) = \frac{1}{\sqrt{2\pi}\sigma_0} e^{-\frac{(x_0 - \bar{x}_0)^2}{2\sigma_0^2}}.$$

2. Furthermore, suppose that before the “game” Player 1 and Player 2 come together and agree that upon receiving his x_0 information, Player 1 will choose his control u_0 to make sure that the state at time $k = 1$ is either $x_1 = \bar{x}_0 + b$ or $x_1 = \bar{x}_0 - b$. Thus, the state x_1 is quantized; the amplitude $b \geq 0$ will jointly be determined by the players before the start of the game. Now, in order to keep the control cost low and thus keep the cost low, the choice of Player 1's control u_0 will be dictated by the distance of x_0 from the two “guideposts” $\bar{x}_0 + b$ and $\bar{x}_0 - b$. In other words, the strategy of Player 1 will be

$$\gamma_0(x_0, \bar{x}_0) = \begin{cases} \bar{x}_0 + b - x_0 & \text{if } x_0 \geq \bar{x}_0, \\ \bar{x}_0 - b - x_0 & \text{if } x_0 < \bar{x}_0. \end{cases} \quad (11)$$

The strategy of Player 1 is illustrated in Fig. 2, where the function

$$f(x_0, \bar{x}_0) \equiv x_1 = x_0 + \gamma_0(x_0, \bar{x}_0)$$




Fig. 2 Strategy of Player 1




Fig. 3 Strategy of Player 1, $\bar{x}_0 = 0$

and in the special case where $\bar{x}_0 = 0$, as in Witsenhausen’s counterexample, the strategy of Player 1 is shown in Fig. 3. The dependence of the strategy of Player 1 on the \bar{x}_0 information is here suppressed.

Since the players come together before the game starts and a cooperative game is played, Player 2 knows the strategy (11) of Player 1. Consequently, at decision time $k = 1$, Player 2 knows that the true state at time $k = 1$ is either $x_1 = \bar{x}_0 + b$ or $x_1 = \bar{x}_0 - b$. Moreover, Player 2 calculates the prior probabilities

$$\mathcal{P}(x_1 = \bar{x}_0 + b) = \mathcal{P}(x_1 = \bar{x}_0 - b) = \frac{1}{2}.$$

Hence, the decision process of Player 2 is now greatly simplified. Player 2 is faced with a binary choice: based on his measurement z_1 , he will have to decide whether the true state x_1 is $\bar{x}_0 + b$ or $\bar{x}_0 - b$ whereupon he will apply his costless control to hopefully set the state $x_2 = 0$ and thus, to the best of his ability, reduce the cumulative cost. Indeed, a typical *communications* scenario is at hand where Player 1 sends one of two possible letters, $\bar{x}_0 + b$ or $\bar{x}_0 - b$, over a Gaussian channel and the job of Player 2 is to *detect* the transmitted letter. To minimize his control effort, Player 1 will decide on which letter to transmit according to its distance from the random initial state x_0 which is known to him.

In view of the strategy (11) employed by Player 1, his expected share of the incurred cost, J_1 , will be

$$\begin{aligned}
J_1 &= K^2 \left[\int_{\bar{x}_0}^{\infty} (\bar{x}_0 + b - x_0)^2 \phi(x_0) dx_0 + \int_{-\infty}^{\bar{x}_0} (\bar{x}_0 - b - x_0)^2 \phi(x_0) dx_0 \right] \\
&= K^2 \left[\int_{\bar{x}_0}^{\infty} (x_0 - \bar{x}_0)^2 \phi(x_0) dx_0 - 2b \int_{\bar{x}_0}^{\infty} (x_0 - \bar{x}_0) \phi(x_0) dx_0 \right. \\
&\quad + b^2 \int_{\bar{x}_0}^{\infty} \phi(x_0) dx_0 + \int_{-\infty}^{\bar{x}_0} (x_0 - \bar{x}_0)^2 \phi(x_0) dx_0 \\
&\quad \left. + 2b \int_{-\infty}^{\bar{x}_0} (x_0 - \bar{x}_0) \phi(x_0) dx_0 + b^2 \int_{-\infty}^{\bar{x}_0} \phi(x_0) dx_0 \right],
\end{aligned}$$

that is,

$$\begin{aligned}
J_1 &= K^2 \left[\int_{-\infty}^{\infty} (x_0 - \bar{x}_0)^2 \phi(x_0) dx_0 + b^2 \int_{-\infty}^{\infty} \phi(x_0) dx_0 \right. \\
&\quad \left. + 2b \left(\int_{-\infty}^0 x \frac{1}{\sqrt{2\pi}\sigma_0} e^{-\frac{x^2}{2\sigma_0^2}} dx - \int_0^{\infty} x \frac{1}{\sqrt{2\pi}\sigma_0} e^{-\frac{x^2}{2\sigma_0^2}} dx \right) \right] \\
&= K^2 \left[\sigma_0^2 + b^2 - 4b \int_0^{\infty} x \frac{1}{\sqrt{2\pi}\sigma_0} e^{-\frac{x^2}{2\sigma_0^2}} dx \right]. \tag{12}
\end{aligned}$$

Now,

$$\int x e^{-\frac{x^2}{2\sigma_0^2}} dx = -\sigma_0^2 e^{-\frac{x^2}{2\sigma_0^2}}$$

and therefore

$$\int_0^{\infty} x e^{-\frac{x^2}{2\sigma_0^2}} dx = \sigma_0^2.$$

Inserting this expression into Eq. (12) yields the cost component

$$J_1(b) = K^2 \left(b^2 - \frac{4}{\sqrt{2\pi}} \sigma_0 b + \sigma_0^2 \right), \quad b \geq 0.$$

The expected contribution of the actions of Player 1 to the cost functional (10), J_1 , is parameterized by his choice of the amplitude/signalling level b . In [1], Witsenhausen chose the amplitude $b = \sigma_0$.

Remark. The expected contribution of Player 1 to the cost functional (10) is minimized when the signalling level $b^* = \sqrt{\frac{2}{\pi}} \sigma_0$, whereupon

$$J_1^* = \left(1 - \frac{2}{\pi} \right) (K\sigma_0)^2. \tag{13}$$

The strategy of Player 2 is the following detection algorithm:

$$\gamma_1(\bar{x}_0, z_1) = \begin{cases} \bar{x}_0 + b & \text{if } z_1 \geq \bar{x}_0, \\ \bar{x}_0 - b & \text{if } z_1 < \bar{x}_0. \end{cases} \quad (14)$$

We now draw the reader's attention to the fact that for Witsenhausen's formulation of the control problem to model a communications scenario and be consistent, one must tacitly assume that before the kickoff the initial state's statistic information \bar{x}_0 is shared with Player 2, as is evident in Eq. (14) above. This, of course, would draw less attention and would appear to be less of an issue if $\bar{x}_0 = 0$, whereupon the strategy of Player 2

$$\gamma_1(z_1) = \begin{cases} b & \text{if } z_1 \geq 0 \\ -b & \text{if } z_1 < 0 \end{cases}$$

would somewhat misleadingly look like Eq. (5).

Concerning the contribution of Player 2 to the cost functional (9), we calculate the probabilities of the possible outcomes:

$$\begin{aligned} \mathcal{P}(x_2 = 0) &= \mathcal{P}(x_0 \geq \bar{x}_0, z_1 \geq \bar{x}_0) + \mathcal{P}(x_0 < \bar{x}_0, z_1 < \bar{x}_0) \\ &= \mathcal{P}(x_0 \geq \bar{x}_0, \bar{x}_0 + b + v_1 \geq \bar{x}_0) + \mathcal{P}(x_0 < \bar{x}_0, \bar{x}_0 - b + v_1 < \bar{x}_0) \\ &= \mathcal{P}(x_0 \geq \bar{x}_0, v_1 \geq -b) + \mathcal{P}(x_0 < \bar{x}_0, v_1 < b) \\ &= \mathcal{P}(x_0 \geq \bar{x}_0)\mathcal{P}(v_1 \geq -b) + \mathcal{P}(x_0 < \bar{x}_0)\mathcal{P}(v_1 < b) \\ &= \frac{1}{2}\mathcal{P}(v_1 \geq -b) + \frac{1}{2}\mathcal{P}(v_1 < b) \\ &= \frac{1}{2}[\mathcal{P}(v_1 \geq -b) + \mathcal{P}(v_1 < b)] \\ &= \frac{1}{2}[\mathcal{P}(v_1 < b) + \mathcal{P}(v_1 < b)] \\ &= \mathcal{P}(v_1 < b) \\ &= \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right], \end{aligned}$$

where

$$\operatorname{erf}(x) \equiv \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt.$$

Now,

$$\bar{x}_0 + b - (\bar{x}_0 - b) = 2b$$

and we calculate

$$\begin{aligned}
 \mathcal{P}(x_2 = 2b) &= \mathcal{P}(x_0 \geq \bar{x}_0, z_1 < \bar{x}_0) \\
 &= \mathcal{P}(x_0 \geq \bar{x}_0, \bar{x}_0 + b + v_1 < \bar{x}_0) \\
 &= \mathcal{P}(x_0 \geq \bar{x}_0, v_1 < -b) \\
 &= \mathcal{P}(x_0 \geq \bar{x}_0) \mathcal{P}(v_1 < -b) \\
 &= \frac{1}{2} \mathcal{P}(v_1 < -b) \\
 &= \frac{1}{4} \left[1 - \operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right].
 \end{aligned}$$

Similarly,

$$\bar{x}_0 - b - (\bar{x}_0 + b) = -2b$$

and we calculate

$$\begin{aligned}
 \mathcal{P}(x_2 = -2b) &= \mathcal{P}(x_0 < \bar{x}_0, z_1 > x_0) \\
 &= \mathcal{P}(x_0 < \bar{x}_0, \bar{x}_0 - b + v_1 > x_0) \\
 &= \mathcal{P}(x_0 < \bar{x}_0, v_1 > b) \\
 &= \mathcal{P}(x_0 < \bar{x}_0) \mathcal{P}(v_1 > b) \\
 &= \frac{1}{2} \mathcal{P}(v_1 > b) \\
 &= \frac{1}{4} \left[1 - \operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right].
 \end{aligned}$$

This allows us to calculate the expected contribution of Player 2 to the cost functional (10),

$$\begin{aligned}
 J_2(b) &= 0 \cdot \mathcal{P}(x_2 = 0) + 4b^2 \cdot \mathcal{P}(x_2 = 2b) + 4b^2 \cdot \mathcal{P}(x_2 = -2b) \\
 &= 2 \left[1 - \operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] b^2.
 \end{aligned}$$

The total cost, $J_1 + J_2$, is

$$J(b) = \left[K^2 + 2 \left(1 - \operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right) \right] b^2 - 2\sqrt{\frac{2}{\pi}} K^2 \sigma_0 b + (K\sigma_0)^2.$$

Without loss of generality assume $\sigma = 1$. Also set

$$b := \frac{1}{\sqrt{2}}b$$

so that as a function of the yet to be determined signalling level b , the cumulative cost (10) is

$$J(b) = 2[K^2 + 2(1 - \text{erf}(b))]b^2 - \frac{4}{\sqrt{\pi}}K^2\sigma_0b + (K\sigma_0)^2. \quad (15)$$

When Player 1 uses a binary signalling level of $\pm b$ —see Eq. (11)—and Player 2 uses the detection strategy (14), the cost function is $J(b)$. It can be further reduced if Player 2 modifies his strategy as follows. Since he cannot be absolutely sure about the correct outcome of the detection step as specified by the strategy (14), Player 2 hedges his bet, does not go all the way, and uses a modified strategy, parameterized by $\alpha \in R^1$, $|\alpha| < 1$:

$$u_1 = \gamma_1(\bar{x}_0, z_1) = \begin{cases} (1 - \alpha)(\bar{x}_0 + b) & \text{if } z_1 \geq \bar{x}_0, \\ (1 - \alpha)(\bar{x}_0 - b) & \text{if } z_1 < \bar{x}_0. \end{cases}$$

As a result, the contribution of Player 2 to the expected cost is

$$\begin{aligned} J_2(\alpha, b; \bar{x}_0) &= \frac{1}{4} \left[1 + \text{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] [\alpha^2(\bar{x}_0 + b)^2 + \alpha^2(\bar{x}_0 - b)^2] \\ &\quad + \frac{1}{4} \left[1 - \text{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] [(\bar{x}_0 + b - (1 - \alpha)(\bar{x}_0 - b))^2 \\ &\quad + (\bar{x}_0 - b - (1 - \alpha)(\bar{x}_0 + b))^2] \\ &= \frac{1}{2} \left[1 + \text{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] (\bar{x}_0^2 + b^2)\alpha^2 \\ &\quad + \frac{1}{2} \left[1 - \text{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] (\bar{x}_0^2 + b^2) \\ &\quad + \frac{1}{2} \left[1 - \text{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] (\bar{x}_0^2 + b^2)(1 - \alpha)^2 \\ &\quad - \left[1 - \text{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] (\bar{x}_0^2 - b^2)(1 - \alpha) \\ &= \frac{1}{2} \left[1 + \text{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] (\bar{x}_0^2 + b^2)\alpha^2 \\ &\quad + \frac{1}{2} \left[1 - \text{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] (\bar{x}_0^2 + b^2)(2 - 2\alpha + \alpha^2) \end{aligned}$$

$$\begin{aligned}
& - \left[1 - \operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] (\bar{x}_0^2 - b^2)(1 - \alpha) \\
& = \alpha^2 \bar{x}_0^2 + (2 - 2\alpha + \alpha^2)b^2 - 2(1 - \alpha)b^2 \operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right).
\end{aligned}$$

The minimum of J_2 is attained when the parameter

$$\alpha^* = \frac{b^2}{\bar{x}_0^2 + b^2} \left[1 - \operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right] \quad (\geq 0)$$

When $\bar{x}_0 = 0$, as in Witsenhausen's counterexample, the optimal parameter

$$\alpha^* = 1 - \operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right)$$

and the expected contribution of Player 2 to the cost is reduced to

$$J_2^*(b) = \left[1 - \operatorname{erf}^2 \left(\frac{b}{\sqrt{2}\sigma} \right) \right] b^2.$$

By optimally hedging his bets, Player 2 has reduced by 50 % his expected contribution to the cost (10). Thus, when $\bar{x}_0 = 0$, as in Witsenhausen's counterexample, the expected cumulative cost $J = J_1 + J_2^*$ is

$$J(b) = (K\sigma_0)^2 + K^2 \left(b^2 - \frac{4}{\sqrt{2\pi}}\sigma_0 b \right) + \left[1 - \left(\operatorname{erf} \left(\frac{b}{\sqrt{2}\sigma} \right) \right)^2 \right] b^2.$$

Without loss of generality, assume $\sigma = 1$. Also set

$$b := \frac{1}{\sqrt{2}}b$$

so that the expected cumulative cost as a function of the amplitude/signalling level b is

$$J(b) = 2[K^2 + 1 - (\operatorname{erf}(b))^2]b^2 - \frac{4}{\sqrt{\pi}}K^2\sigma_0 b + (K\sigma_0)^2, \quad b > 0. \quad (16)$$

Assume the parameter $\sigma_0 \gg 1$. The cumulative cost (10) then attains a local minimum at $b^* \approx \frac{1}{\sqrt{\pi}}\sigma_0$ and consequently the value of the functional (10) is

$$J^* \approx \left(1 - \frac{2}{\pi} \right) (K\sigma_0)^2. \quad (17)$$




Fig. 4 Expected cost as a function of signalling level

This is also the minimal contribution of Player 1 to the expected cost—see Eq. (13). Hence, we conclude that when $\sigma_0 \gg 1$ and the binary signalling communications protocol is used, the global minimum is given by Eq. (13) and it is attained when Player 1 uses the amplitude/signalling level

$$b^* = \sqrt{\frac{2}{\pi}}\sigma_0.$$

We will focus on the benchmark/canonical scenario where the parameters $K = 0.2$ and $\sigma_0 = 5 (> 1)$. The cumulative expected cost function (16) is depicted in Fig. 4. The optimal binary signalling level is

$$b^* = 2.874875$$

and the expected minimal cumulative cost

$$J^* = 0.3651961.$$

As expected, and by design, in this scenario, the contribution of Player 2 to the cumulative cost is small. From Fig. 4 it is also plainly evident that the cost function is not convex and it has a local minimum.

Binary signalling was also used in Witsenhausen's original counterexample [1] where the signalling level is

$$b = \sigma_0$$

but in [1] the strategy $\gamma_1^*(z_1)$ of Player 2 did not entail a detection protocol and instead a *continuous* function of his measurement z_1 is used. The expected cumulative cost in Witsenhausen's paper is

$$J_W^* = 0.404253$$

and we see that

$$J_W^* > J^* = 0.3651961.$$

Players 1 and 2 could agree on a multilevel signalling/communications protocol—a staircase-like rendition of a multilevel signalling protocol [6, 7], is shown in Fig. 5. This allows Player 1 to choose the signalling level which is closest to the randomly selected initial state x_0 , thereby reducing his control effort. At the same time, the use of multiple signalling levels complicates the detection task of Player 2 and the probability of him erring increases. This, in turn, increases the expected contribution




Fig. 5 Multilevel threshold strategy of Player 1




Fig. 6 Multilevel graded strategy of Player 1

of Player 2 to the cumulative cost and a trade-off is needed. A $3\frac{1}{2}$ -level signalling protocol, as in [3], reduces the expected cost to $J^* = 0.1673132$.

A piecewise constant staircase signalling protocol—see Fig. 5—is not optimal and the horizontal rungs should be slightly slanted upward. The cost can be further reduced if Player 1 uses a continuous signalling protocol as illustrated in Fig. 6.

The best result so far, using a continuous, monotonically increasing, signalling function and a continuous “detection” algorithm, is reported in [4]. In [4] the parameters are $K = 0.5, \sigma_0 = 10$, and $K = 0.25, \sigma_0 = 10$. When our suboptimal binary signalling protocol is employed and, as in [4], the “non-canonical” parameters are $K = 0.5$ and $\sigma_0 = 10$, the cost function $J(b)$ is shown in Fig. 7. The optimal binary signalling level is then

$$b^* = 5.645646$$

and the expected minimal cumulative cost

$$J^* = 9.084513.$$

When, as in [4], the problem parameters are $K = 0.25, \sigma_0 = 10$ and our suboptimal binary signalling protocol is employed, the cost function $J(b)$ is shown in Fig. 8. The optimal signalling level is

$$b^* = 5.645646$$

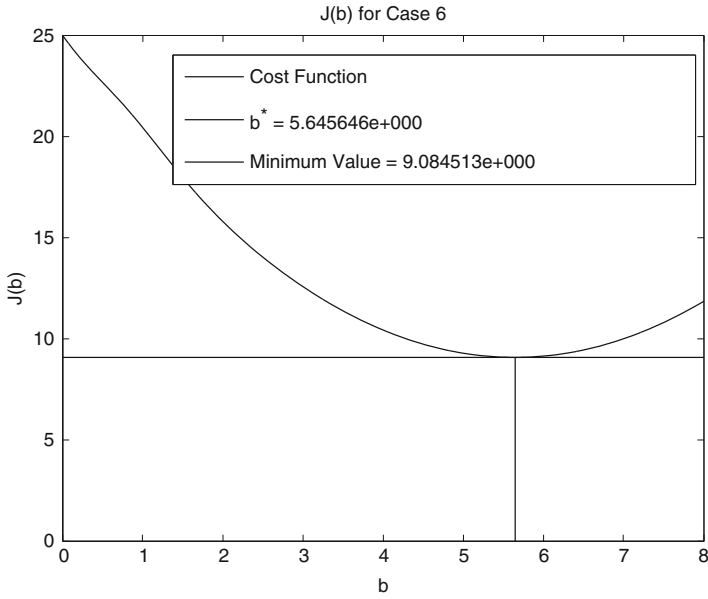


Fig. 7 Expected cost as a function of signalling level




Fig. 8 Expected cost as a function of signalling level

and the expected minimal cumulative cost

$$J^* = 2.271128.$$

In summary, in decentralized control, the players are not supposed to communicate, and yet in Witsenhausen's communication and control scenario, the players share the initial state's statistic \bar{x}_0 and, moreover, are tacitly allowed to communicate before the kickoff of the cooperative "game" and establish a communication protocol—one thus refers to signalling. Indeed, in Witsenhausen's original counterexample, two letters are transmitted over a Gaussian channel, that is, the signalling strategy of Player 1 is

$$\gamma_0(x_0) = \begin{cases} \sigma_0 & \text{if } x_0 \geq 0, \\ -\sigma_0 & \text{if } x_0 < 0. \end{cases}$$

The job of Player 2, the receiver, is reduced to detection; however in Witsenhausen's counterexample the conventional thresholding—based detection scheme is not used, and instead an optimal continuous decision function akin to a squashing function is used. We have shown that using a detection strategy reduces the cost compared to Witsenhausen's continuous decision function. The reader is referred to Appendix A where the basics of detection theory are outlined.

Thus, while not explicitly presented as such, Witsenhausen's counterexample centers on the design of a suboptimal communications protocol. Note however that communication before the kickoff of the game would not be needed and the optimal control problem would be a bit more decentralized if the players could independently arrive at their respective optimal strategies.

4 No Communication During the Game

Since communicating/signalling over a noise-corrupted channel is so integral to Witsenhausen's decentralized control problem suboptimal solutions, it is instructive to take things to an extreme and consider the following decentralized optimal control problem where during run time no communication is to take place: the decentralized control problem is specified by Eqs. (1), (2), (4), (6) and the initial state information is specified in (1), or (7), as before, but now, during runtime, Player 2 receives no information whatsoever, that is, at time $k = 1$ a measurement of the state x_1 is *not* taken by Player 2. In other words, the extreme case is now considered where the parameter $\sigma \rightarrow \infty$. It is also assumed that over time the random initial state x_0 of the dynamical system presented to Player 1 will be chosen according to the probability distributions (1) or (7), and this is known to both Players 1 and 2. Thus, the respective strategies of Players 1 and 2 are of the form (9) and

$$u_1 = \gamma_1(\bar{x}_0).$$

We intentionally now take a “wrong turn” and initially analyze this special case from the vantage point of the widely used communications paradigm discussed in Sect. 3, where the players come together and establish a communication protocol before kickoff time.

4.1 *The Players Communicate Before the Game*

Although Player 2 operates in an open-loop mode and no signalling will take place, not everything is lost: since a cooperative control scenario is considered where the players are allowed to communicate *prior* to the start of the game, they could as well agree that Player 1 will always see to it that, irrespective of the realized initial state x_0 , at decision time $k = 1$ the state $x_1 = b$, always; the optimal amplitude b is yet to be determined. Player 1 is now committed to an affine strategy

$$\gamma_0(x_0, \bar{x}_0) = b - x_0 \quad (18)$$

and Player 2 knows this. Hence, Player 2 is absolved of even taking a measurement of x_1 —he does not need the measurement z_1 and he will always apply the control

$$u_1 = b,$$

driving the state x_2 to 0.

The amplitude b must be decided on ahead of time. In order to minimize his average cost, Player 1 will calculate the optimal amplitude b^* by solving the optimization problem

$$J_1^* = \min_b E_{x_0}(K^2(b - x_0)^2)$$

which yields

$$b^* = \bar{x}_0$$

so that the optimal strategy of Player 1 is affine and is

$$\gamma_0^*(x_0, \bar{x}_0) = \bar{x}_0 - x_0. \quad (19)$$

Hence, the realized minimal cost is

$$J_1^*(x_0) = K^2(x_0 - \bar{x}_0)^2 \quad (20)$$

and consequently, the average minimal cost of Player 1 is,

$$J_1^* = K^2\sigma_0^2. \quad (21)$$

The minimal expected cost (21) is *not* dependent on the initial state's statistic \bar{x}_0 . Also,

$$J_2^* = J_1^*.$$

Note: Similar to the stochastic games paradigm, it is assumed throughout that the σ_0 statistic is shared by the players.

Indeed, Player 2 does not need to know the initial state's statistic \bar{x}_0 : Since Player 1 communicates the amplitude b^* to Player 2 before the start of the game, then, similar to Player 1, upon analyzing the optimization problem at hand, Player 2 will independently arrive at the conclusion that the initial state's statistic is in fact $\bar{x}_0 = b^*$ and he will apply the control

$$u_1^* = \bar{x}_0. \quad (22)$$

Example Assume that prior communication is allowed and $\sigma \rightarrow \infty$. When the problem parameters are the canonical parameters $K = 0.2$ and $\sigma_0 = 5$, the “minimal” average cost J_1^* and the expected “minimal” cost J_2^* are

$$J_1^* = J_2^* = 1.$$

In particular, if $\bar{x}_0 = 0$ then the “optimal” strategies are $\gamma_0^*(x_0) = -x_0$ and $u_1^* = 0$.

In conclusion, the strategy (19) of Player 1 and the open-loop optimal control (22) of Player 2 are commensurate with the herein stipulated information pattern, irrespective of whether Player 2 is privy to the initial state's statistic \bar{x}_0 . For the “minimal” average cost (21) to be realizable when prior communication is allowed, Player 2 does not need to know the initial state's statistic \bar{x}_0 . This is so because prior communication takes place and the “game” is cooperative. Indeed, prior communication allows Player 2 to infer the initial state's statistic \bar{x}_0 on his own.

However, should the initial state's statistic \bar{x}_0 be known to both Players 1 and 2, and if, in addition, the strategy (19) of Player 1 and the open-loop control (22) of Player 2 were indeed optimal, that is, $K^2\sigma_0^2$ is *the* minimal expected cost, then no prior communication would be required: based on the public information \bar{x}_0 available to them, both players would independently solve the decentralized optimal control problem, arrive at their respective optimal strategies, and calculate their expected costs—this being predicated on the assumption that Player 1 is after minimizing his average cost. Strictly speaking, a unique Nash equilibrium would have been obtained. In this instance, and courtesy of the solution of the optimal control problem, implicit communication would automatically materialize. It however turns out that the modulation strategy (19) of Player 1 which was derived using the communications model from Sect. 3 is *not* optimal, as will be shown in the next section where the optimal solution is derived.

4.2 *The Players Do Not Communicate Before the Game*

The information pattern is s.t. Player 2 knows that Player 1 knows that his measurement of the initial state x_0 is \bar{x}_0 , according to the distribution (7), or, alternatively, Player 1 knows that the initial state x_0 presented to him will be drawn from the distribution (7) and Player 2 knows this as well. Thus, the strategy of Player 1 is given by Eq. (9) and the strategy of Player 2 is

$$u_1 = \gamma_1(\bar{x}_0).$$

Since the statistic \bar{x}_0 is public information and it is not a random variable, the strategies' dependence on \bar{x}_0 is temporarily suppressed and we shall refer to the strategy

$$u_0 = \gamma_0(x_0)$$

of Player 1 and the control

$$u_1 \in R^1$$

of Player 2.

Concerning the correct analysis of the optimization process:

1. First, take the point of view of Player 1, who has been provided the initial state information x_0 : Player 1 is playing against the optimal input $u_1^* \in R^1$ of Player 2, and since his strategy is a delayed commitment strategy, no random variables feature in his optimization. Thus, his cost function

$$J^{(1)}(u_1^*) \equiv \min_{u_0 \in R^1} [K^2 u_0^2 + (x_0 + u_0 - u_1^*)^2]$$

and consequently his optimal control must satisfy the relationship

$$u_0^* = -\frac{1}{K^2 + 1}(x_0 - u_1^*).$$

Hence, the optimal strategy of Player 1 and the optimal control of Player 2 must satisfy the relationship

$$\gamma_0^*(x_0) = -\frac{1}{K^2 + 1}(x_0 - u_1^*). \quad (23)$$

2. Next, take the point of view of Player 2, who is playing against the optimal strategy $\gamma_0^*(x_0)$ of Player 1 and as far as he is concerned, the initial state is a random variable whose p.d.f. is specified by Eq. (7). Thus, his cost functional

$$J^{(2)}(\gamma_0^*(\cdot)) \equiv \min_{u_1 \in R^1} \{E_{x_0} (K^2 (\gamma_0^*(x_0))^2 + [x_0 + \gamma_0^*(x_0) - u_1]^2 | \bar{x}_0)\}$$

and consequently his optimal control must satisfy the relationship

$$u_1^* = E_{x_0}(x_0 + \gamma_0^*(x_0) \mid \bar{x}_0)$$

that is,

$$u_1^* = \bar{x}_0 + E_{x_0}(\gamma_0^*(x_0) \mid \bar{x}_0). \quad (24)$$

Both Players 1 and 2 calculate the expectations of the L.H.S. and R.H.S. of Eq. (23) and obtain the equation

$$E_{x_0}(\gamma_0^*(x_0) \mid \bar{x}_0) = -\frac{1}{K^2 + 1}(\bar{x}_0 - u_1^*). \quad (25)$$

Inserting Eq. (25) into Eq. (24) yields

$$u_1^* = \bar{x}_0. \quad (26)$$

In other words, the optimal strategy of Player 2 is

$$\gamma_1^*(\bar{x}_0) = \bar{x}_0 \quad (27)$$

and inserting Eq. (26) into Eq. (23) yields the optimal strategy of Player 1

$$\gamma_0^*(x_0, \bar{x}_0) = \frac{1}{K^2 + 1}(\bar{x}_0 - x_0). \quad (28)$$

Having obtained the optimal strategies, the respective value functions of Players 1 and 2 are calculated as follows:

$$V_0^{(1)}(x_0, \bar{x}_0) = \frac{1}{K^2 + 1}(x_0 - \bar{x}_0)^2 \quad (29)$$

and

$$V_0^{(2)}(\bar{x}_0) = \frac{1}{K^2 + 1}\sigma_0^2. \quad (30)$$

The analysis from above is summarized in

Theorem 1. *The special case of Witsenhausen's decentralized optimal control problem (1), (2), (4), and (6), where the parameter $\sigma \rightarrow \infty$, but with the slightly more general initial state information specified by Eq. (7), is considered. Thus, the case is considered where at time $k = 1$ a measurement of the state x_1 is not taken by Player 2. The respective optimal strategies of Players 1 and 2 are linear and are given by Eqs. (28) and (27) and their value functions are given by Eqs. (29) and (30).*

Remark. The value of Player 2 is equal to the average cost/value of Player 1, that is,

$$V_0^{(2)}(\bar{x}_0) = E_{x_0}(V_0^{(1)}(x_0, \bar{x}_0) | \bar{x}_0). \quad (31)$$

Corollary 2. *In the special case where, as in Witsenhausen's paper, the initial state's statistic $\bar{x}_0 = 0$, the optimal strategies are*

$$\gamma_0^*(x_0) = -\frac{1}{1 + K^2}x_0, \quad (32)$$

and

$$\gamma_1^* = 0 \quad (33)$$

and the players' value functions are

$$V_0^{(1)}(x_0) = \frac{1}{1 + K^2}x_0^2 \quad (34)$$

and

$$V_0^{(2)} = \frac{1}{1 + K^2}\sigma_0^2. \quad (35)$$

When the parameter (the canonical parameter)

$$K = 0.2$$

the slope of the linear control law of Player 1 is ≈ -1 and the function

$$\begin{aligned} f^*(x_0) &\equiv x_0 + \gamma_0^*(x_0) \\ &= \frac{K^2}{1 + K^2}x_0 \\ &\approx 0.04x_0. \end{aligned}$$

In summary, in our analysis the somewhat unconventional information pattern germane to Witsenhausen's problem where at time $k = 0$ the information about the initial state's statistic \bar{x}_0 is shared among the players and which therefore somewhat detracts from its perceived degree of decentralization, has been retained. It is however important to realize that in this special case of Witsenhausen's problem, where the variance of the measurement error of Player 2 is very large, the *optimal* solution has been obtained. As a result, no additional communication among the players before kickoff time is needed in order to establish a communications protocol, as is tacitly assumed in the case where the parameter σ is finite, the optimization problem is much harder and has not yet been solved, and one must

fall back on suboptimal solutions derived using the communications/signalling paradigm discussed in Sect. 3. Since in the special case treated herein where the parameter $\sigma \rightarrow \infty$ the optimal solution can be independently obtained by both Players 1 and 2, *no* communication prior to kickoff is needed and therefore one is entitled to say that the attendant optimal control problem has been solved in a more decentralized manner.

An additional instance where an *optimal* solution can be easily obtained entails the symmetric information pattern where Player 1 does not have access to the initial state information and both players' information is shared and is specified by Eq. (7); in addition, Player 1 knows that Player 2 has the public information (7) and, vice versa, Player 2 knows that Player 1 has the public information (7). In this case, the strategy of Player 1 is the control $u_0 \in R^1$ and the strategy of Player 2 is the control $u_1 \in R^1$.

1. First, take the point of view of Player 1: Player 1 is playing against the optimal input $u_1^* \in R^1$ of Player 2 and now the random variable x_0 features in his optimization. Thus, his cost function

$$J^{(1)}(u_1^*) \equiv \min_{u_0 \in R^1} E_{x_0} ([K^2 u_0^2 + (x_0 + u_0 - u_1^*)^2])$$

and consequently his optimal control must satisfy the relationship

$$u_0^* = -\frac{1}{K^2 + 1}(\bar{x}_0 - u_1^*).$$

2. Next, take the point of view of Player 2, who is playing against the optimal strategy/control u_0^* of Player 1 and as far as he is concerned, the initial state is a random variable whose p.d.f. is specified by Eq. (7). Thus, his cost functional

$$J^{(2)}(u_0^*) \equiv \min_{u_1 \in R^1} \{ E_{x_0} (K^2 (u_0^*)^2 + [x_0 + u_0^* - u_1]^2 \mid \bar{x}_0) \}$$

and consequently his optimal control must satisfy the relationship

$$u_1^* = \bar{x}_0 + u_0^*. \quad (36)$$

This yields the optimal controls/strategies

$$u_0^* = 0, \quad (37)$$

$$u_1^* = \bar{x}_0 \quad (38)$$

and the optimal/minimal expected cost

$$J^* = \sigma_0^2 \quad (39)$$

The analysis from about is summarized in

Theorem 2. *The special case of Witsenhausen's decentralized optimal control problem (1), (2), (4), and (6), where the parameter $\sigma \rightarrow \infty$, but with the slightly more general initial state information specified by Eq. (7), is considered. Thus, the case is now considered where at time $k = 1$ a measurement of the state x_1 is not taken by Player 2 and, in addition, Player 1 does not have access to the initial state information x_0 . The respective optimal strategies of Players 1 and 2 are given by Eqs. (36) and (37) and the minimal cost is given by Eq. (38). The minimal cost is not dependent on the initial state's statistic \bar{x}_0 .*

5 Conclusion

In Witsenhausen's problem statement the following must be made clear. The decision problem is only partially decentralized. At time $k = 0$ the information on the initial state's statistic \bar{x}_0 is exchanged among the players. In addition, the synthesis of suboptimal solutions rests on the assumption that before the game starts, during "foreplay," the players are allowed to come together and establish a communication protocol. This entails allowing Player 1 to convey the information on the initial state's statistic to Player 2. Thus, a communications problem using a Gaussian communications channel is modeled. Now, communication can be referred to as signalling, although, in the informational economics literature [9] the term *signalling* assumes a somewhat different meaning. Alternatively, it is tacitly assumed that at decision time $k = 0$, Player 2 takes a measurement of the initial state and communicates his measurement to Player 1. In this case Player 1 knows that Player 2 thinks that the initial state is distributed according to Eqs. (1) or (7). Evidently, the control problem is not completely decentralized and the strategy of Player 1, which naturally incorporates all the information available to him at decision time $k = 0$, has the somewhat unconventional form (9). This state of affairs is masked if, as in Witsenhausen's problem statement, it is assumed that the statistic $\bar{x}_0 = 0$.

The problem with Witsenhausen's problem goes beyond the somewhat hidden requirement that the initial state's statistics information be shared by Players 1 and 2, which immediately detracts from the decentralized aspect of the control problem: the suboptimal solutions are based on the perception that a cooperative communication problem is at hand and this requires the Players to come together and agree on a communications protocol prior to the kickoff of the game. Obviously, the better the communications protocol is, the lower will be the expected cost, and so, when viewing Witsenhausen's problem in the context of a mathematical model of a communications scenario, suboptimal solution methods readily suggest themselves. Now, the fact that the initial state's statistics information is shared is perhaps OK, but the additional requirement that the players come together before the kickoff of the game and agree on a communications protocol, as is indeed the case

in the classical communications paradigm is giving one pause for thought. No such thing would be required if the *optimal* solution of Witsenhausen’s problem would be known, in which case the two intelligent Players 1 and 2 could independently figure out their respective modulation and detection strategies. This unfortunately is not the case, because the *optimal* solution of Witsenhausen’s problem is not yet fully known and therefore the game is not playable without the artificial preliminary step of setting up a (suboptimal) communications protocol. Knowledge of the optimal solution would obviate the need for this preliminary step. There is one exception: in the special case investigated in Sect. 4 where the variance of the measurement error of Player 2 is very big, the optimal strategies are linear and are known—we refer to the players’ optimal strategies (23) and (27); the point is that both players can *independently* derive their optimal strategies.

In summary, Witsenhausen’s problem is not fully decentralized to start with and in the absence of an optimal solution the players must establish a somewhat artificial communication protocol before kickoff time. As such, Witsenhausen’s problem is somewhat contrived.

Appendix A: Witsenhausen’s Counterexample and Detection Theory

An information theoretic analysis of Witsenhausen’s binary signalling protocol follows. The case where a system could be in either one of two states is considered.

Specifically, when the binary signalling protocol is invoked in Sects. 3 and 4, the state could be $x_1 = b$ or $x_1 = -b$. Suppose the true state x_1 is not known; however a measurement of the state x_1 is taken, whereupon, based on the measurement’s outcome, a declaration concerning the state x_1 is made, namely, x_1 is declared to be b or, alternatively, the state x_1 is declared to be $-b$. Let the performance of the said classifier be quantified by the receiver operating characteristic (ROC) which is a *confusion matrix*

True/Rep.	$x_1 = b$	$x_1 = -b$
b	P_{TR}	$1 - P_{\text{FTR}}$
$-b$	$1 - P_{\text{TR}}$	P_{FTR}

Here, the probabilities

$$P_{\text{TR}} \equiv \mathcal{P}(x_1 = b \mid b)$$

and

$$P_{\text{FTR}} \equiv \mathcal{P}(x_1 = -b \mid -b).$$

Concerning the detection task, consider the typical situation where prior information on the state x_1 of the system is available and, based on a received measurement of the state x_1 , one is about to decide whether the state $x_1 = b$ or $x_1 = -b$. The performance of the classifier, that is, its ROC, is quantified by the above specified confusion matrix. The following holds.

Theorem. *The state x_1 is known to be either $x_1 = b$ or $x_1 = -b$ and the prior information on x_1 is*

$$\mathcal{P}(x_1 = b) = p.$$

Suppose a measurement is about to be taken whereupon, based on the recorded measurement, a classifier will either declare the state $x_1 = b$ or, alternatively, the state $x_1 = -b$. The performance of the classifier is characterized by its confusion matrix which is parametrized by P_{TR} and P_{FTR} . Then the expected information gain will be

$$\begin{aligned} I(P_{\text{TR}}, P_{\text{FTR}}, p) &= pP_{\text{TR}} \log \left(\frac{P_{\text{TR}}}{pP_{\text{TR}} + (1-p)(1-P_{\text{FTR}})} \right) \\ &\quad + p(1-P_{\text{TR}}) \log \left(\frac{1-P_{\text{TR}}}{p(1-P_{\text{TR}}) + (1-p)P_{\text{FTR}}} \right) \\ &\quad + (1-p)(1-P_{\text{FTR}}) \log \left(\frac{1-P_{\text{FTR}}}{pP_{\text{TR}} + (1-p)(1-P_{\text{FTR}})} \right) \\ &\quad + (1-p)P_{\text{FTR}} \log \left(\frac{P_{\text{FTR}}}{p(1-P_{\text{TR}}) + (1-p)P_{\text{FTR}}} \right). \end{aligned} \quad (40)$$

The expected information gain function (40) is depicted in Fig. 9.

Concerning the classification algorithm proper, the probabilities P_{TR} and P_{FTR} , which characterize the classifier's performance, are obtained as follows.

The performance of the sensor used to measure the state x_1 is specified by Eq. (41): if the classifier's threshold is set to t , the detection algorithm yields the estimate \hat{x}_1 of the state x_1 according to

$$\hat{x}_1 = \begin{cases} b & \text{if } z_1 \geq t, \\ -b & \text{if } z_1 < t. \end{cases} \quad (41)$$

Hence, from Fig. 10—see, e.g., [8]—one concludes that the detection probability

$$P_{\text{TR}} \equiv \mathcal{P}(\hat{x}_1 = b \mid x_1 = b)$$

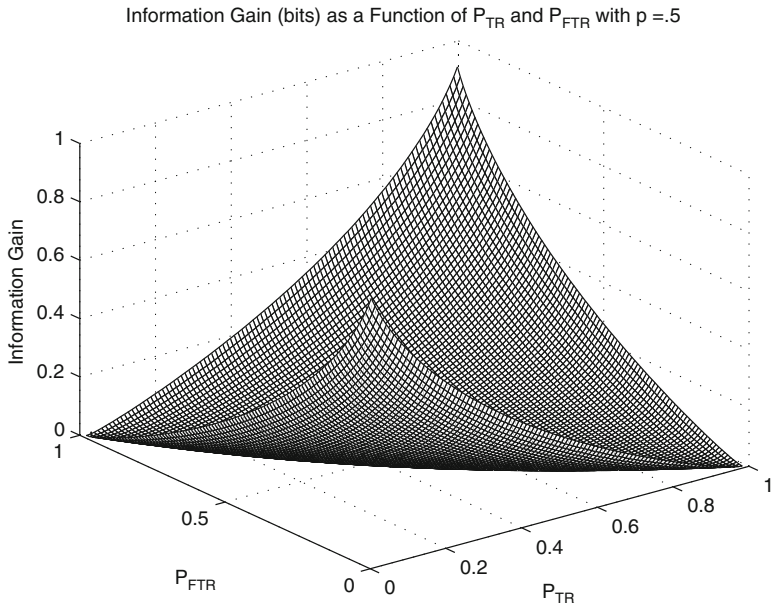


Fig. 9 Expected information gain




Fig. 10 Calculation of the probabilities of detection and false alarm

$$\begin{aligned}
 &= \mathcal{P}(z_1 \geq t) \\
 &= 1 - \phi\left(\frac{t - b}{\sigma}\right).
 \end{aligned}$$

Similarly, with reference to Fig. 10, the “False Alarm” probability P_{FA} is calculated as follows:

$$P_{FA} \equiv \mathcal{P}(\hat{x}_1 = b \mid x_1 = -b)$$

$$\begin{aligned}
&= \frac{1}{\sqrt{2\pi}\sigma} \int_t^\infty e^{-\frac{(x-(-b))^2}{2\sigma^2}} dx \\
&= \frac{1}{\sqrt{2\pi}\sigma} \int_t^\infty e^{-\frac{(x+b)^2}{2\sigma^2}} dx \\
&= \frac{1}{\sqrt{2\pi}\sigma} \int_{t+b}^\infty e^{-\frac{y^2}{2\sigma^2}} dy \\
&= 1 - \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{t+b} e^{-\frac{y^2}{2\sigma^2}} dy
\end{aligned}$$

that is,

$$P_{\text{FA}} = 1 - \phi\left(\frac{t+b}{\sigma}\right).$$

Hence,

$$\begin{aligned}
P_{\text{FTR}} &= 1 - P_{\text{FA}} \\
&= \phi\left(\frac{t+b}{\sigma}\right).
\end{aligned}$$

Now,

$$\phi(x) = \frac{1}{2} \left[1 + \operatorname{erf}\left(\frac{x}{\sqrt{2}}\right) \right].$$

Using the signal to noise ratio (SNR) definition

$$\text{SNR} \equiv \frac{b}{\sigma}$$

and non-dimensionalizing the threshold

$$t := \frac{t}{\sigma}$$

we finally obtain

$$P_{\text{TR}}(t, \text{SNR}) = \frac{1}{2} \left[1 - \operatorname{erf}\left(\frac{t - \text{SNR}}{\sqrt{2}}\right) \right] \quad (42)$$

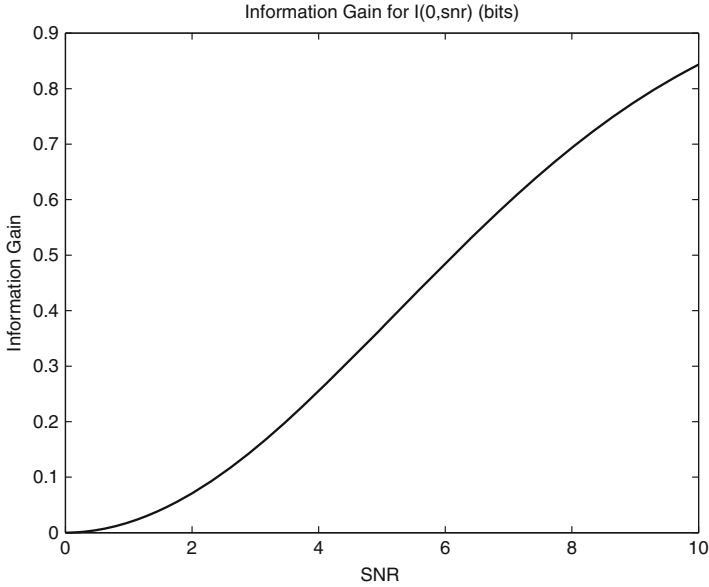


Fig. 11 Expected information gain as a function of SNR

and

$$P_{\text{FTR}}(t, \text{SNR}) = \frac{1}{2} \left[1 + \text{erf} \left(\frac{t + \text{SNR}}{\sqrt{2}} \right) \right]. \tag{43}$$

Due to symmetry, the prior probability that the state $x_1 = b$ is $p = \frac{1}{2}$. Inserting the expressions (42) and (43) into Eq. (40) with $p = \frac{1}{2}$ directly yields the expected information gain as a function of the threshold setting t and the SNR. Due to symmetry, the threshold will be set to $t = 0$. The expected information gain as a function of the SNR is shown in Fig. 11. For the optimal b^* calculated in Sect. 3 the attendant SNR yields the expected maximal information gain: it is 1 bit.

References

1. Witsenhausen, H.S.: A counterexample in stochastic optimal control. *SIAM J. Control* **6**(1), 131–147 (1968)
2. Li, N., Marden, J.R., Shamma, J.S.: Learning approaches to the witsenhausen counterexample from a view of potential games. In: *Joint 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference, Shanghai, PR China*, pp. 157–162, 16–18 Dec 2009
3. Lee, J.T., Lau, E., Ho, Y.-C.: The Witsenhausen counterexample: a hierarchical search approach for non-convex optimization problems. *IEEE Trans. Automatic Control* **46**(3), 382–397 (2001)

4. McEneaney, W.M., Han, S.H., Liu, A.: An optimization approach to the Witsenhausen counterexample. In: IEEE Conference on Decision and Control, Orlando, FL, Dec 2011
5. Basar, T.: Variations on the theme of the Witsenhausen counterexample. In: Proceedings of the 47th IEEE Conference on Decision and Control Cancun, Mexico, pp. 1614–1619, 9–11 Dec 2008
6. Park, S.Y., Grover, P., Sahai, A.: A constant-factor approximately optimal solution to the Witsenhausen Counterexample. In: Joint 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference, Shanghai, P.R. China, pp. 2882–2886, 16–18 Dec 2009
7. Sahai, A., Grover, P.: Demystifying the Witsenhausen counterexample. *IEEE Control Syst. Mag.* **30**, 20–24 (2010)
8. Selin, I.: *Detection Theory*. Academic, New York (1965)
9. Spence, M.: Job market signaling. *Q. J. Econ.* **87**(3), 355–374 (1973)