An analytic solution for Sannikov's "optimality equation" in the case of a single binding incentive constraint

Asaf Plan

November 16, 2011

This note presents an analytic solution for Sannikov's (2007) "optimality equation" in a special case. Sannikov introduces and analyzes a class of continuous-time games with imperfect public monitoring. He introduces the optimality equation, an ordinary differential equation that describes the boundary of the set of equilibrium payoffs, $\mathcal{E}$. Based on the optimality equation, Sannikov presents techniques for computing $\mathcal{E}$, to be carried out with the assistance of a computer. Sannikov mentions that "it would be interesting to explore other computational procedures to find the set $\mathcal{E}$." In the general case, the optimality equation is nonlinear and no analytic solution is apparent. This note presents an analytic solution in the special case where only one player at a time has a binding incentive constraint. That condition holds in the restrictive case where only one player at a time departs from the static best response. It may also hold more broadly under an alternative monitoring structure where there is a single signal which identifies the difference between the two players' actions. We derive closed-form comparative statics for the largest, symmetric equilibrium payoff in a simple game.

Consider the following three stage games,

P.D.

	0	1
0	$0,\ 0$	$\beta,\ -\chi$
1	$-\chi,\ \beta$	$\beta-\chi,\ \beta-\chi$

Shifted P.D.

	0	1
0	$\chi,\ \chi$	$\beta+\chi,\ 0$
1	$0,\ \beta+\chi$	$\beta,\ \beta$

"Give or Take"

	0	1
0	$\chi,\ \chi$	$\beta+\chi,\ 0$
1	$0,\ \beta+\chi$	$0,\ 0$

where $\beta > \chi > 0$. On the left is a standard prisoner's dilemma. (Action 1 is "cooperate" while 0 is "defect".) Sannikov presents an example of a partnership game with $\beta = 2$ and $\chi = 1$. In the middle, all payoffs have simply been shifted upward by χ, which will be convenient below. On the right, payoffs for joint cooperation have been reduced to zero. The rightmost game, which I call "give or take," is simpler because in extremal equilibria only one player cooperates at a time.

Recall the continuous-time signal structure from Sannikov's partnership example: For each player $i = 1, 2$, there is an independent signal, Ψ^i, which evolves according to the stochastic differential equation (SDE),

$$d\Psi_t^i = A_t^i dt + dB_t^i, \tag{1}$$

where A_t^i is player i's action at time t and B^1, B^2 are a pair of independent, standard Brownian motions. Cooperation by player i can be enforced by positively conditioning her continuation payoffs on Ψ^i.

In either the second or third stage game above, suppose that at time t player one is cooperating (playing action 1) while two is defecting (playing 0). Player one could increase her flow payoff by an amount χ by deviating. In an extremal equilibrium, we enforce this action profile by conditioning the evolution of player one's payoff on her signal as follows,

$$dW_t^1 = rW_t^1 dt + r\chi\left(d\Psi_t^1 - 1dt\right) = \underbrace{rW_t^1}_{=\mu_t^1} dt + \underbrace{r\chi}_{=\sigma_t^1} dB_t^1. \tag{2}$$

Recall that the drift, μ_t^1, is determined by the "promise-keeping" condition: Because player one's current flow payoffs are zero while her expected discounted average payoff is W_t, we must have positive drift $\mu_t^1 = rW_t^1$. The volatility, σ_t^1, is the smallest value sufficient to incentivize player one to cooperate. If the volatility were smaller, then we could not be in an equilibrium, as player one would prefer to deviate. On the other hand, if the volatility were greater, we could not be in an extremal equilibrium, because incentives are costly. Suppose that the upper half of the boundary of the set of equilibrium payoffs can be described by a function f where $W_t^2 = f(W_t^1)$ during times where player one cooperates while two defects. Player two's payoffs evolve as follows,

$$dW_t^2 = f'(W_t^1)dW_t^1 = f'(W_t^1)rW_t^1 dt + \underbrace{f'(W_t^1)r\chi}_{=\sigma_t^2} dB_t^1.$$

Notice $\sigma_t^2 = f'(W_t^1)\sigma_t^1$.

What I have presented so far is a small bit of Sannikov's analysis. The main observation underlying this note is that when only one player's incentive constraint is binding, the payoff process of that player is a known and somewhat tractable stochastic process — an Ornstein-Uhlenbeck (O-U) process of the "exploding" type.[1] This is true for player one above, as seen in equation 2. Player two's payoffs do not follow such a nice process, but we can solve for W^2 in terms of W^1 along this part of the boundary of $\mathcal{E}$, which is a central result in this note. (The relationship $W_t^2 = f(W_t^1)$ is one way to describe the boundary of $\mathcal{E}$. Sannikov's optimality equation instead regards the support function of $\mathcal{E}$.)

In contrast, consider the case where at time t, both players are cooperating in the middle game under the signal structure described above. Here each player's incentive constraint is binding with respect to her own signal, and each player's payoff also varies indirectly with the other player's signal through the relationship $W_t^2 = f(W_t^1)$. The SDEs are

$$\begin{aligned} dW_t^1 &= r(W_t^1 - \beta)dt + r\chi dB_t^1 + \frac{1}{f'(W_t^1)dt} r\chi dB_t^2 \\ dW_t^2 &= r(W_t^1 - \beta)dt + f'(W_t^1)r\chi dB_t^1 + r\chi dB_t^2. \end{aligned} \tag{3}$$

The processes W^1 and W^2 are interdependent, neither is an O-U process, and I am not aware that anything can be said about the boundary of $\mathcal{E}$ beyond Sannikov's optimality equation. One can again write down an ODE for $W_t^2 = f(W_t^1)$, but in this case that ODE is nonlinear, and no analytic solution is apparent.

In the Give or Take game only one player's incentive constraint binds at a time because only one player cooperates at a time, while the other plays the static best response, defection. Such equilibria, which exclude simultaneous cooperation, are of limited interest. However, equilibria with simultaneous cooperation may become similarly tractable under certain monitoring structures where there is only a single, one-dimensional

[1] Recall that an Ornstein-Uhlenbeck process takes the form $dX_t = -\lambda(X_t - m)dt + \sigma dB_t$. It is more common to consider values $\lambda > 0$, in which case the process is mean-reverting. Here we have $\lambda = -r < 0$. This process drifts away from zero, and the further away it is, the greater the drift.

signal. In place of the two independent signals Ψ^1 and Ψ^2 described above, suppose there is only one signal:

$$d\Psi_t^0 = A_t^1 dt - A_t^2 dt + dB_t^0, \tag{4}$$

here dB_t^0 is a one-dimensional, standard Brownian motion.[2] (Any monitoring structure where there is only a single, one-dimensional signal fails to meet Sannikov's Assumption 1, though I believe that for this particular signal structure together with any of three stage games mentioned before, all of Sannikov's conclusions still hold as if Assumption 1 was satisfied.[3]) In the give or take game, this one signal would yield the same equilibrium payoff set as the two signals described before; our analytic solution holds in either case. In the prisoner's dilemma, the one and two signal cases are qualitatively different. With the one signal, only a single player's incentive constraint binds even while both are cooperating, thus our analytic solution holds.

1 The solution for the "Give or Take" game

Recall the stage game,

"Give or Take"

	0	1
0	$\chi,\ \chi$	$\gamma,\ 0$
1	$0,\ \gamma$	$0,\ 0$

where I have substituted $\gamma = \beta + \chi$, where $\gamma > 2\chi$.

In degenerate cases, the set of equilibrium payoffs $\mathcal{E}$ contains only the static Nash payoff, (χ, χ). Otherwise, $\mathcal{E}$ looks something like a convex tear drop that is symmetric about the 45 degree line and has its point at the static Nash payoff. Here $\mathcal{E}$ is the same whether we consider the original monitoring structure (1) or the single-signal monitoring structure (4). Clearly $\mathcal{E}$ is homogeneous of degree one in (γ, χ), so one might normalize $\chi = 1$, but I do not as I think doing so may obscure more than it simplifies.

Suppose we are on the portion of the boundary of $\mathcal{E}$ above the 45 degree line and player one's current expected discounted average payoff is $W_t^1 = v$. I will solve for player two's current expected discounted average payoff, $W_t^2 = f(v)$. The current action profile is $(1, 0)$; player one is "giving", which requires incentives, while two is "taking," which is the static best response. The evolution of v is described by equation 2; it has drift rv and volatility $r\chi$. The drift of $f(v)$, that is $\mathrm{E}[df(v)]/dt$, is given by Ito's lemma. We then have the Hamilton-Jacobi-Bellman equation,

$$f(v) = \gamma + \frac{\mathrm{E}[df(v)]}{r\,dt} = \gamma + vf'(v) + \frac{r\chi^2}{2} f''(v),$$

which is a linear ODE. It has the general solution,

$$f(v) = \gamma + K_1 v - K_2 \left(\frac{1}{2} \exp\left(- \left(\frac{v}{\chi\sqrt{r}} \right)^2 \right) + \frac{v}{\chi\sqrt{r}} \int_0^{v/(\chi\sqrt{r})} e^{-s^2} ds \right), \tag{5}$$

[2] This is a limiting case of the two-signal structure: $d\Psi_t^i = A_t^i dt + \frac{1}{\sqrt{2}} dB_t^i + \theta dB_t^3$, as $\theta \to \infty$, where B^1, B^2, B^3 are independent, standard Brownian motions.

[3] This signal structure does not meet Sannikov's Assumption 1, that any two action profiles can be statistically distinguished from one another. In this case, the action profile $(1, 1)$ cannot be distinguished from $(0, 0)$, but doing so is unnecessary. Assumption 1 implies that deviations of different players can be statistically distinguished. Under the monitoring structure proposed here we cannot distinguish the deviation from $(1, 0)$ to $(0, 0)$ from the deviation to $(1, 1)$, but the latter deviation is not a concern. We can still distinguish the deviation from $(1, 1)$ to $(1, 0)$ from the deviation to $(0, 1)$, which is necessary for the optimality equation to hold.

where K_1 and K_2 are constants to be determined from the boundary conditions. (The corresponding ODE that arises when both players' incentive constraints are binding is not linear, and I cannot write down any solution.[4]) The solution for f above is not closed-form, as it contains the integral $\int e^{-s^2} ds$ which famously has no closed-form solution.[5]

The boundary conditions follow from symmetry: $f(\chi) = \chi$, $f(\bar{v}) = \bar{v}$ and $f'(\bar{v}) = -1$ for some $\bar{v}$ that we must solve for. The point $(\bar{v}, \bar{v})$ is the largest, symmetric payoff profile in $\mathcal{E}$; we know $\bar{v} \in [\chi, \gamma/2)$. The third boundary condition relies both on symmetry and the fact that the boundary of $\mathcal{E}$ is differentiable everywhere except at (χ, χ).

There are three boundary conditions and three unknowns: $K_1, K_2, \bar{v}$. Using the third boundary condition to pin down K_1 yields

$$f(v) = \gamma - v - K_2 \underbrace{\left(\frac{1}{2} \exp\left(- \left(\frac{v}{\chi\sqrt{r}} \right)^2 \right) - \frac{v}{\chi\sqrt{r}} \int_{v/(\chi\sqrt{r})}^{\bar{v}/(\chi\sqrt{r})} e^{-s^2} ds \right)}_{=\phi(v)}$$

(notice the limits of integration have changed). An understanding of the solution will require an understanding of the term in the large parenthesis, $\phi(v)$, which I will return to. For now, note that $\phi > 0$, $\phi' \leq 0$ and $\phi'' < 0$ for $v \in [\chi, \bar{v}]$. Either of the remaining boundary conditions can be used to pin down K_2, and the last one then pins down $\bar{v}$:

$$\frac{\gamma - 2\chi}{\phi(\chi)} = K_2 = \frac{\gamma - 2\bar{v}}{\phi(\bar{v})}. \tag{6}$$

The latter equality for K_2 is simpler because the integral disappears in $\phi(\bar{v})$, but for the moment the former may offer an easier interpretation: $f(v) = \gamma - v - (\gamma - 2\chi)\phi(v)/\phi(\chi)$, which can be rearranged as,

$$\frac{\gamma - (v + f(v))}{\gamma - 2\chi} = \frac{\phi(v)}{\phi(\chi)}, \tag{7}$$

which equals one at $v = \chi$ and is decreasing in v. On the LHS, the numerator is the amount lost to inefficiency, and the denominator is the joint, static gain from cooperation. As Sannikov describes, the path of continuation payoffs (W_t^1, W_t^2) will eventually be absorbed at the static Nash payoff, (χ, χ). The ratio on the LHS of (7) is the expected discounted share of time following such Nash absorption. This is a normalized measure of inefficiency. As $r \to 0$ (that is, the player's become very patient), the expected discounted share of time following Nash absorption goes to 0, no matter the initial value of $v > \chi$.

[4] For example, given equations (3), we have,

$$W^2 = f(v) = \beta + vf'(v) + \frac{r\chi^2}{2} \left(1 + \frac{1}{(f'(v))^2} \right) f''(v).$$

where $v = W^1 - \beta$. Notice the term $f''/(f')^2$.

[5] That integral is the "error function," which is related to the cumulative distribution function of the normal distribution.

1.1 Comparative statics for $\bar{v}$

Recall, the point $(\bar{v}, \bar{v})$ is the largest, symmetric payoff profile in $\mathcal{E}$. The value $\bar{v}$ is the largest solution in the range $[\chi, \gamma/2)$ for equation 6, which I rearrange:

$$0 = g(\bar{v}) = \frac{\gamma - 2\chi}{\gamma - 2\bar{v}} - \underbrace{\left(\frac{1}{2}\exp\left(-\left(\frac{1}{\sqrt{r}}\right)^2\right) - \frac{1}{\sqrt{r}}\int_{1/\sqrt{r}}^{\bar{v}/(\chi\sqrt{r})} e^{-s^2}\,ds\right)}_{=\phi(\chi)}\underbrace{\left(2\exp\left(\left(\frac{\bar{v}}{\chi\sqrt{r}}\right)^2\right)\right)}_{=1/\phi(\bar{v})}. \tag{8}$$

There is a degenerate solution: $\bar{v} = \chi$. As we discuss in the next section, that is the only solution if the players are impatient (large r) and/or cooperation yields little benefit ($(\gamma - 2\chi)/\chi$ small), in which case $\mathcal{E}$ contains only the static Nash payoff point, (χ, χ). For now suppose that players are patient and/or cooperation yields great benefit, so $\bar{v} > \chi$. We can implicitly differentiate (8) to derive comparative statics for $\bar{v}$ with regards to γ and r. While $g(\bar{v})$ is not itself closed-form, the resulting comparative statics are closed-form in terms of $\bar{v}$, which is nice.

Differentiating both sides of $0 = g(\bar{v}(\gamma), \gamma)$ yields,

$$\begin{aligned}
\frac{d\bar{v}}{d\gamma} &= -\frac{\frac{\partial g}{\partial \gamma}}{\frac{\partial g}{\partial \bar{v}}} \\
&= -\frac{-\frac{2(\bar{v}-\chi)}{(\gamma-2\bar{v})^2}}{\left[\frac{2}{\gamma-2\bar{v}}\left(\frac{\gamma-2\chi}{\gamma-2\bar{v}}\right)\right] + \left[\frac{2\bar{v}}{\chi^2 r}\frac{\phi(\chi)}{\phi(\bar{v})} - \frac{2\chi}{\chi^2 r}\right]} \\
&= \frac{\chi^2 r(\bar{v} - \chi)}{\chi^2 r(\gamma - 2\chi) + \gamma\left(\bar{v} - \chi\right)\left(\gamma - 2\bar{v}\right)},
\end{aligned}$$

which is positive, as was to be expected. Notice this expression contains $\bar{v}$, but no ϕ terms.[6]

The value $\bar{v}$ is increasing in γ in a mechanical way. Perhaps a more meaningful measure of the extent of cooperation is the expected discounted share of time for which the player's cooperate (before eventually being absorbed at the static Nash equilibrium), call this m. When the player's payoffs are at $(\bar{v}, \bar{v})$, this value solves $\bar{v} = (1 - \bar{l})\chi + \bar{l}\gamma/2$. So we have,

$$\begin{aligned}
\bar{l} &= \frac{2(\bar{v} - \chi)}{\gamma - 2\chi} \\
\frac{d\bar{l}}{d\gamma} &= \frac{2\bar{v}'(\gamma)(\gamma - 2\chi) + 2(\bar{v} - \chi)}{(\gamma - 2\chi)^2} \\
&= \bar{l}\frac{1}{(\gamma - 2\chi)}\left(1 + \frac{1}{\left(1 + \frac{1}{2r}\frac{\gamma}{\chi}\frac{\gamma-2\chi}{\chi}\bar{l}(1 - \bar{l})\right)}\right),
\end{aligned}$$

which is again increasing, as was to be expected.

We can similarly derive $\frac{d\bar{v}}{dr}$.

[6] The derivative $g'(\bar{v})$ contains a term $\frac{\phi(\chi)}{\phi(\bar{v})}$. We have eliminated this between the second and third step here, by replacing $\frac{\phi(\chi)}{\phi(\bar{v})} = \frac{\gamma-2\chi}{\gamma-2\bar{v}}$, which follows from $g(\bar{v}) = 0$.

1.2 When is any cooperation possible?

As we mentioned in the previous section, the value $\bar{v}$ is largest solution in the range $[\chi, \gamma/2)$ of equation (8). I am interested in the range of parameters such that some cooperation is possible, that is $\bar{v} > \chi$:

$$\Theta = \left\{ \left(r, \frac{\gamma - 2\chi}{\chi} \right) : \ \exists \bar{v} \in (\chi, \gamma/2), \ g(\bar{v}) = 0 \right\}.$$

The value $g(v)$ is monotone decreasing in γ and increasing in r, so for some increasing function $h(\cdot)$,

$$\Theta = \left\{ \left(r, \frac{\gamma - 2\chi}{\chi} \right) : \ 0 < r \leq h\left(\frac{\gamma - 2\chi}{\chi} \right) \right\},$$

as expected.

Over the relevant range, $(\chi, \gamma/2)$, the function $g(v)$ varies smoothly, and it attains positive values as it approaches either boundary. The question is whether it dips below zero over some intermediate range. If we are on the boundary of Θ, then it must be that g is just tangent to zero: $g'(\bar{v}) = 0$ and $g(\bar{v}) = 0$. At such a point,

$$\frac{\chi^2 r}{\bar{v}} \frac{1}{\gamma - 2\bar{v}} \left(\frac{\gamma - 2\chi}{\gamma - 2\bar{v}} \right) + \frac{\chi}{\bar{v}} = \frac{\phi(\chi)}{\phi(\bar{v})} = \frac{\gamma - 2\chi}{\gamma - 2\bar{v}}.$$

Those two equalities characterize the boundary of Θ.

1.3 Relation to an Ornstein-Uhlenbeck hitting-time problem

While player one cooperates, her continuation payoffs evolve as a particular Ornstein-Uhlenbeck (abbreviated OU) process of the exploding type. In turn, player two's continuation payoffs are determined by the expected discounted hitting times and splitting probabilities of that OU process. Calculating such hitting times and splitting probabilities is a common type of exercise in probability theory. Here our interest is in determining the upper bound $\bar{v}$, which solves a particular hitting-time problem.

Consider the Ornstein-Uhlenbeck process that is absorbed at lower boundχ, reflected at upper bound $\bar{v}$, and in between obeys the SDE

$$dV_t = rV_t dt + r\chi dB_t$$

where B is a standard Brownian motion. (This process represents the continuation payoff of the player that is presently cooperating, $V_t = \min\{W_t^1, W_t^2\}$.) Let τ be the first hitting time, $V_\tau = \chi$, that is the absorption time. Given the initial value $V_0 = v$, consider the expected discounted absorption time,

$$l(v; r, \chi, \gamma, \bar{v}) = \mathrm{E}[re^{-r\tau} | V_0 = v].$$

That is r times the Laplace transform of τ with parameter $r > 0$. Solving for the function l is a common type of problem in probability theory.[7]

The function l depends on the parameters r, χ and $\bar{v}$. In our context, the reflecting barrier $\bar{v}$ solves

$$l(\bar{v}; r, \chi, \gamma, \bar{v}) = \frac{2(\bar{v} - \chi)}{\gamma - 2\chi}.$$

[7] In fact, probabilists are more interested in solving for the undiscounted expectation of the absorption time, $\mathrm{E}[\tau | V_0 = v]$. A standard way to do so is solving for the Laplace transform, which is easier, and then inverting it, which is often not so easy. Conveniently here the object of interest is the Laplace transform itself.

This ties $\bar{v}$ to a standard probability exercise. We showed before that $l(\bar{v}) = 1 - \phi(\bar{v})/\phi(\chi)$, where ϕ is defined in equation (8).

2 Prisoner's dilemma with a single signal

Consider the stage game

Shifted P.D.

	0	1
0	$\chi,\ \chi$	$\beta+\chi,\ 0$
1	$0,\ \beta+\chi$	$\beta,\ \beta$

where $\beta > \chi > 0$, along with the alternative monitoring structure with the one signal

$$d\Psi_t^0 = A_t^1 dt - A_t^2 dt + dB_t^0, \tag{9}$$

where dB_t^0 is a one-dimensional, standard Brownian motion.

Under this signal structure, the incentive constraints only bind on one player at a time, even while both players are cooperating. While just player one is cooperating, her payoff process is the same as in equation 2 above. While both players are cooperating there are two, symmetric cases to consider. If $f'(W_t^1) \equiv \frac{dW_t^2}{dW_t^1} \leq -1$, then the incentive constraint is only binding on player one, and payoffs follow the SDEs,

$$\begin{aligned} dW_t^1 &= r(W_t^1 - \beta)dt + r\chi dB_t^0 \\ dW_t^2 &= r(W_t^1 - \beta)dt + f'(W_t^1) r\chi dB_t^0. \end{aligned}$$

Again W^1 is a tractable O-U process (now shifted upward by β). Here we get the incentives for player two for free.[8] Symmetrically, if both players are cooperating and $f'(W_t^1) \geq -1$, then W^2 is an O-U process, while we get the incentives for player one for free.

We can solve for $\mathcal{E}$ in terms of two complementary portions of its upper boundary. Over the first portion, player one plays 1 while two plays 0, as in Give or Take. This portion is above the 45-degree line, begins at $w^1 = w^2 = \chi$ and continues to some point $w^1 = \hat{v}$ that will be described momentarily. Here $w^2 = f(w^1)$ where f has the same general solution as in (5). Over the second portion, both players play 1, but the incentive constraint binds only on player one. This portion begins at $w^1 = w^2 = \bar{v} > \hat{v}$ and continues to $w^1 = f(\hat{v})$, $w^2 = \hat{v}$. Over this portion $w^2 = f(w^1 - \beta)$, where f again has the same general solution as before.

We have four boundary conditions, of which the first three follow from symmetry as before: $f(\chi) = \chi$, $f(\bar{v} - \beta) = \bar{v}$ and $f'(\bar{v} - \beta) = -1$. The fourth boundary condition describes the optimal transition between the action profiles $(1, 0)$ and $(1, 1)$, which occurs where player one has payoff $\hat{v}$. That condition is

$$f(\hat{v}) = (\beta + \chi) + \hat{v} f'(\hat{v}) + \frac{r\chi^2}{2} f''(\hat{v}) = \beta + (\hat{v} - \beta) f'(\hat{v}) + \frac{r\chi^2}{2\,(f'(\hat{v}))^2} f''(\hat{v}).$$

Together these four boundary conditions and the general solution for f identify the boundary of $\mathcal{E}$. However,

[8] In order to have both players cooperate, we need $\sigma^1 \geq r\chi$ and $\sigma^2 \leq -r\chi$. Here only the first constraint is binding; we have $\sigma^2 = f'(W_t^1)\sigma^1 = f'(W_t^1) r\chi \leq -r\chi$.

the situation is more complicated here than in the previous Give or Take game, both because there is an additional boundary condition and because that boundary condition is more complicated.

3 Conclusion

The optimality equation does not in general yield an analytical solution. I do not view that as a weakness of the model. In the natural sciences many (most?) important problems do not have analytic solutions. Instead, a problem would be considered "solved" upon providing the relevant differential equation and proving the existence and uniqueness of the solution to that equation. That is what Sannikov (2007) has done regarding the set of equilibrium payoffs $\mathcal{E}$ in his class of two-player, continuous-time games of imperfect public monitoring. Nevertheless, it is advantageous to have an analytic solution where possible.

This note provides such an analytic solution for the boundary of $\mathcal{E}$ in the special case where only one player's incentive constraint binds at a time. That case holds when only one player at a time deviates from the static best response, as in the Give or Take game presented in section (1). That case may also hold with simultaneous cooperation under an alternative monitoring structure, as in section (2).

From our analytic solution we derive closed-form comparative statics for the largest, symmetric equilibrium payoff of the Give or Take game.

We relate the solution to the hitting times and splitting probabilities of an Ornstein-Uhlenbeck process of the exploding type. Hitting times and splitting probabilities are standard objects of study in probability theory.

Another natural case where only one player's incentive constraint binds is the setting of principal-agent problems, as in Sannikov (2008). However, our solution relies on the discreteness of the action space, and our solution is simplest in the case of a binary action space. The assumption of a binary action space for the principal is uncommon. (In that case the principal can

References

Sannikov, Y. (2007): "Games with Imperfectly Observable Actions in Continuous Time," *Econometrica*, 75, 1285–1329.

——— (2008): "A Continuous-Time Version of the Principal Agent Problem," *The Review of Economic Studies*, 75, 957–984(28).