Journal of Gambling Issues
Issue 40, December 2018
http://igi.camh.net/doi/pdf/10.4309/jgi.2018.40.7
DOI: http://dx.doi.org/10.4309/jgi.2018.40.7

policy paper

Optimal Public Policy for Government-Operated Gambling

Irwin Lipnowski[1] & Austin McWhirter[1]

[1] Department of Economics, University of Manitoba, Winnipeg, Manitoba, Canada

Abstract

This paper presents a framework for formulating the optimal public policy for government-operated gambling. The goal of public policy with respect to government-operated gambling is typically stated as "harm minimization." This claim masks the possible trade-off between an increase in social harm (H) and the government's incremental net revenue (R) from increased gambling activity. Using a graphical approach, we depict first the feasible combinations of H and R, and then identify the combinations that could be classified as efficient, thereby allowing the minimum social harm for any given level of the government's incremental net revenue from gambling. We indicate how the optimal combination of H and R could be identified and realized in both the short and long run. We then utilize the body of research on gambling and its effects to qualify what this trade-off operates in the real world.

Keywords: trade-off, efficient, harm minimization, disordered gambling, optimal policy

Résumé

Ce document présente un cadre pour l'élaboration d'une politique publique optimale encadrant les jeux de hasard gérés par le gouvernement. L'objectif d'une politique publique en matière de jeu géré par le gouvernement vise généralement une « minimisation des méfaits ». Cela permet de dissimuler le compromis possible entre une augmentation des dommages sociaux (H) et le revenu net supplémentaire du gouvernement (R) provenant d'une augmentation des activités de jeu. À l'aide d'une approche graphique, nous décrivons d'abord les combinaisons possibles entre les méfaits sociaux et les revenus (H et R), puis nous identifions les combinaisons « efficaces », ce qui permet de réduire au minimum les méfaits pour la société à tout niveau des revenus nets supplémentaires tirés du jeu par le gouvernement. Nous indiquons comment la combinaison optimale méfaits/revenus pourrait être établie et réalisée à court et à long terme. Nous utilisons ensuite le

corpus de recherches sur le jeu et ses effets pour qualifier à quoi ressemble ce compromis dans le monde réel.

Introduction

In a world of abundance, the academic discipline of economics would not exist. Scarcity imposes limits on the feasible choices we can make. Acquiring more of one good necessitates foregoing another. A central challenge that society faces is to make the best, or *optimal* use of our scarce resources.

In examining the role of government as a monopoly provider of gambling services, clients who take part in moderate amounts of gambling for recreation and entertainment can be viewed no differently than customers in any ordinary business. Society would not maintain an ethical concern about the marketing and promotion of gambling activity for this segment of its clientele. If gambling promotion could target recreational gamblers exclusively, governments could act like private corporations in pursuit of maximum profit.

Because it is impossible to target recreational gamblers in such a manner when gambling activity is promoted, governments are faced with a trade-off. Although marketing gambling might well allow a government to realize additional net revenue, such marketing would further harm those disordered gamblers who are induced to increase their gambling activity, and it may well increase the prevalence in the population of those who gamble as well. For example, Harrah's, a casino and hotel chain located in several American cities, used marketing to "promote more frequent and longer gambling sessions among avid and experienced gamblers" and managed to double its profits in one year (Quinn, 2001, p. 134). Put simply, as revenue rises, so too does the harm caused to disordered gamblers.

Cast in this stark way, the decision by a government to forego net revenue to reduce the social harm to disordered gamblers and their families would appear to be morally compelling. When additional revenue for a government is placed on one side of a balance scale while a composite of social harms—an increase in the unsustainable losses of disordered gamblers, a potential increase in criminal activity to support their addiction, the additional stress and depression suffered by them and their families, an increase in domestic violence and family breakdowns, and the real possibility that their state of hopelessness could lead them to suicide—are placed on the other side of the balance scale, the appropriate choice for society seems obvious.

However, the conflict between money, in the form of additional tax revenue, versus the increase in social harm as two competing options, has been presented in the gambling research literature in a decidedly biased way, exemplified through the contributions of Dickerson (2003) and Jackson, Christensen, Francis, and Dowling (2015).

The funds that are generated through the operation of casinos and lotteries can be used to reduce harm in society at large. For example, additional tax revenue could save lives by funding needle exchange programs, or by supplying Naloxone kits to neutralize the fatal effects of fentanyl. It could also be used to fund after-school programs for low-income children or many other socially worthwhile programs. It is plausible to believe that the harm caused by government-sponsored gambling could be offset by the benefits to society from the use of the resulting profits.[1] Indeed, when one weighs the harm created by the government-operated gambling activity against the harm reduction that could, potentially, be realized if the government were to deploy effectively the proceeds from its gambling operations, the choice between opposing or supporting the government's involvement in gambling on ethical grounds is far from obvious.

Professional economists generally regard their discipline as a value-free social science that provides politicians with a set of tools for evaluating alternative policies. The social goals to which these tools are applied are treated by economists as exogenously given by the politicians—who, in a highly idealized and vastly over-simplified world, are deemed to be expressing a social consensus, or "the will of the people". In other words, the relative value of the benefits that flow from the government's net revenue and the costs to society of the harm caused by disordered gambling are ultimately subjective, and are determined by the preferences and opinions of society. Perhaps the electorate believes that if the money from the casinos can pay for 150 more hospital beds at the cost of potentially ruining the life of one disordered gambler, then it is a worthwhile trade for society. Perhaps the lives saved from these additional hospital beds are worth more than the damage caused to disordered gamblers and their families. Providing definitive guidance to answer such questions is neither the domain of the economist nor of this paper.

The far more modest technical task undertaken by economists is to determine all of the feasible alternative outcomes that are attainable by society and then to identify the optimal outcome, based on exogenously given preferences of society.

In this paper, we present a conceptual framework to analyze the trade-off between the government's incremental net revenue from its gambling operations and the resulting harm, and we illustrate how, in principle, one could determine the optimal

[1]Blaszczynski, Sharpe & Walker (2001) appear to have diluted the term "harm minimisation." For example, in their final report, *The Assessment of the Impact of the Reconfiguration on Electronic Gaming Machines as Harm Minimisation Strategies for Problem Gambling*, they assert that "harm minimisation" is "used synonymously with harm reduction" (Blaszczynski et al., p. 23, 2001). Such usage provides little guidance to a policymaker facing many difficult competing choices. Harm reduction is simply one of a multitude of social *desiderata*, such as improvements for the general population in literacy, numeracy, nutrition and longevity.

combination of the two variables, given society's preferences, both in the short and long run. We then illustrate how this framework might be applied in the real world, using selective evidence from the vast research literature on gambling.

Uncovering the Trade-Off: A Look at Efficiency

To simplify the exposition, we will describe outcomes in terms of only two variables: an index of social harms attributable to gambling activity (H) and the incremental net revenue from government-operated gambling activity (R).

Before we begin, we must first define what we mean by the government's incremental net revenue. For a governmental monopoly firm operating a gambling enterprise, its net revenue (or profit) would be the amount by which its total revenue exceeds its total cost. Its total revenue would be the total amount spent by customers on gambling activities. Its total costs would be the sum of the payments to gamblers, the rent paid (or imputed) for the use of its facilities, a depreciation charge for its buildings and equipment (if it owns these), money spent on the repair and maintenance of its facilities, and expenditures on all operating costs (including marketing and promotion, heating, lighting, air quality control, security, administrative and employee salaries). Tax payments, if any, could be treated simply as a transfer from one pocket of the government to another.

Unfortunately, this picture of government's net revenue (or profit) significantly overstates the net contribution that a government's gambling operations can make to funding hospitals, education, infrastructure, and so on. As Henriksson and Lipsey (1999, p. 260) have observed, the government's monopoly profit from gambling fails to take into account the taxes that are foregone from other sources with the introduction of government-operated gambling activities. For example, the introduction of a casino into a community may well divert customers from other businesses in the area, since money that could have been spent at existing establishments is instead spent at the casino. Following the introduction of casinos in Atlantic City in 1978, one-hundred of the city's two-hundred and fifty restaurants were driven out of business (Henriksson, 1996, p. 117). As a result of this "crowding out effect," the government lost tax revenue that it could have levied if those businesses had remained open or performed better. The government's incremental net revenue that results from its introduction of gambling is its monopoly profit from gambling activities minus the tax revenue that is foregone from the crowding out effect of the gambling activities (Lipsey and Henriksson, 1999, p. 260). Therefore, this paper defines R as the government's incremental net revenue to reflect the actual benefits realized by society from the presence of government-operated gambling.

Having clearly defined the variables we will be using, we can proceed to develop our framework to shed light on the trade-off between H and R. Consider the B_0 curve in Figure 1. All points on or above the B_0 curve, such as points A, B and C, are feasible; that is, given the possible ways the government can manage its gambling operations, these combinations of H and R are attainable.

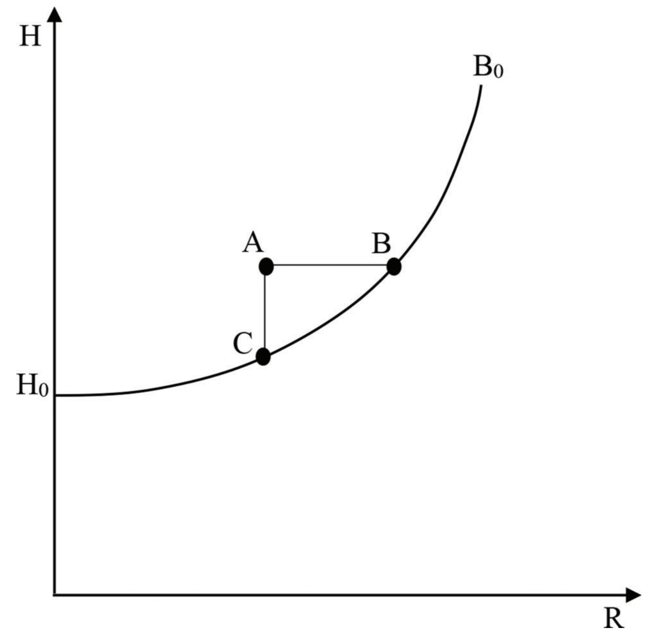

Figure 1. The B_0 curve depicts the efficient boundary of feasible R and H combinations. Above the B_0 curve R can be increased without increasing H and H can be reduced without decreasing R. Points below B_0 are not attainable.

An example of a change in the way in which gambling operations are managed, and thus a movement from one feasible point to another, would-be casinos reducing their hours of operation. Such action has been taken by some countries as a way of curbing disordered gambling (Quinn, 2001, p. 136). Presumably, H would be reduced, but so too would R, since the casino would forego revenue it could have realized had its hours of operation not been curtailed. This change could be imagined as moving from B to C in Figure 1, since there is a trade-off between H and R. More generally, reducing both H and R could be represented by moving from any feasible point to another that is to the left and below the original point.

Although all of the points on or above the B_0 curve are feasible, only those points on the B_0 curve are efficient. Efficiency is defined here as a combination of H and R in which the only way to reduce H is to also reduce R, and the only way to increase R is, at the same time, to increase H. Although point A is feasible, it is clearly inefficient, since we could move to point C and achieve less harm with the same incremental net revenue, or to point B and achieve more incremental net revenue while holding the harm from disordered gambling constant. A movement from point A to any point that is strictly between B and C on the B_0 curve would achieve greater incremental net revenue and a reduction in harm. A similar improved outcome would be possible from any feasible point that is not on the B_0 curve. Since optimality can only be found at an efficient point, the optimal combination of R and H must accordingly lie on the B_0 curve.

At any feasible point above the B_0 curve, someone could be made better off without making anyone else worse off. A move from A to C would allow disordered gamblers and their families to become better off while allowing the beneficiaries of government programs that are funded by the incremental net revenue from gambling to become no worse off.[2] Likewise, the beneficiaries of government expenditures could become better off if the government were to move from A to B, while disordered gamblers would be no worse off than before the move. But, one cannot move from any point on the B_0 curve to any other feasible point without making someone worse off. Moving from one efficient point to another involves a trade-off between harm and incremental net revenue. Efficient points along the B_0 curve have the property that no individual can be made better off without making someone else worse off.

The general shape of the B_0 curve reflects the assumption that each additional dollar of the government's incremental net revenue from gambling results in social harm that increases at an increasing rate. If harm increased at a constant rate for each additional dollar of incremental net revenue, B_0 would be a straight line. Note that even if the government were to abandon completely gambling as a source of revenue, the social harm from gambling activity could not be lowered to a level below H_0. The presence of privately operated gambling sites online or gambling venues abroad, for example, would still result in a level of social harm denoted H_0 in Figure 1. Thus, only levels of harm that exceed H_0 are attributable to the government's involvement in gambling operations.

The Quest for the Optimal Efficient Combination of H and R

Once one accepts that good social policy should ensure that society be somewhere on the B_0 curve, the question remains "Where along the B_0 curve does the socially optimal combination of the government's incremental net revenue and social harm occur?"

To answer this question, one must know society's preferences. For example, how many family breakdowns because of disordered gambling is society willing to accept to fund one additional after-school program from gambling revenue? If society could choose between a combination of three family breakdowns and four additional after-school programs or two family breakdowns and three after-school programs, which would it choose? In the general form, we must know how society ranks all the feasible combinations of H and R.

This idea is elucidated by means of "Social Indifference Curves." Two such curves, denoted I_1 and I_2, are depicted in Figure 2 below. To know which combinations of H and R are preferred by society to other combinations, we must first know which combinations are deemed to be equally preferred. A Social Indifference Curve shows every combination of H and R between which society is "indifferent," assuming that

[2]Indeed, if the beneficiaries of programs funded by gambling revenue care about pathological gamblers, a move from A to C would in fact benefit them.

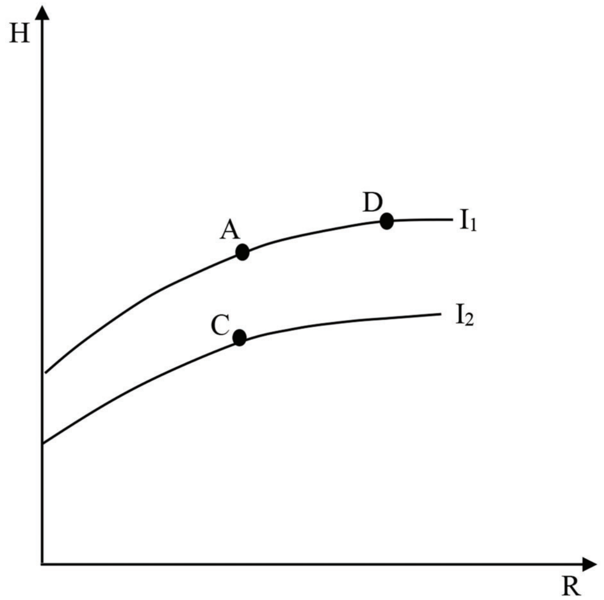

Figure 2. Two Social Indifference Curves, I_1 and I_2; combinations of R and H along I_2 are all equally preferred by society, and all of these are strictly preferred to combinations of R and H along I_1.

society has reached an "ethical consensus" (Mishan, 1980, p. 148). For example, in Figure 2, I_1 tells us that the combination of H and R represented by point A provides the same utility (or happiness) to society as the combination of H and R represented by point D.

Every point on I_2 is strictly preferred by society to every point on I_1. To explain this, consider points A and C. Precisely the same level of R is achieved at both points, but point C also results in less H. Therefore, C is strictly preferred by society to A. Since A is preferred equally to every other point along I_1, C must be preferred to every point along I_1. So, even though points D and C do not appear to be comparable since at D, an increase in R entails greater H, using Social Indifference Curves enables an analyst to confirm that C is preferred to D. One might imagine that D and C represent the two alternative combinations of family breakdowns and after-school programs presented earlier. Only by knowing society's preferences, expressed through Social Indifference Curves, can an appropriate judgment about the superiority of C over D be made.

We assume that a unique Social Indifference Curve will pass through every combination of H and R. In Figure 2 we have, for illustrative purposes, depicted only two of the infinite number of non-intersecting Social Indifference Curves. Thus any combination of H and R can be compared to any other combination of H and R in the entire space of feasible combinations of H and R.

The most preferred feasible combination of H and R can be determined by superimposing Figure 2 onto Figure 1.

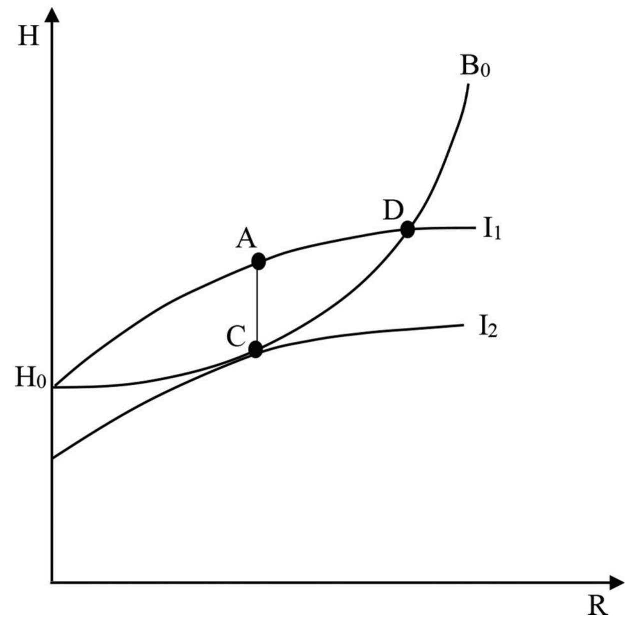

Figure 3. Point C is society's most preferred attainable combination of R and H.

As one can see from Figure 3, an efficient feasible combination of H and R along the B_0 curve that is most preferred by society occurs at point C, where the highest attainable Social Indifference Curve I_2 is tangent to the B_0 curve. In this way, I_2 is the optimal Social Indifference Curve that can be realized (given that all points below the B_0 curve are unattainable). If the government chose any other combination of H and R along the B_0 curve, that point would lie along a less favourable Social Indifference Curve, and society would then be worse off. This is illustrated by point D, which like C is an efficient point, but is less preferred than C because it is on a less preferred Social Indifference Curve. Although every point along the B_0 curve represents an efficient combination of H and R, society's optimal choice is at C, where the maximum social utility, corresponding to Social Indifference Curve I_2, would occur.

The shape of the Social Indifference Curves is not necessarily concave to the origin as depicted in Figures 2 and 3. Their shape is determined by the relative weights that society places on the increase in social harm that results from an increase in the government's incremental net revenue. Figure 4 illustrates the shape of Social Indifference Curves if society held the view that the government is corrupt and would squander every dollar of net revenue generated by gambling activity. In this case, an increase in R would be of no value to society. Social Indifference Curves would then be perfectly horizontal as only H would matter to society. Thus, a lower H is always preferred to a higher H, irrespective of the level of R.

In the highly improbable case in which society had an entirely callous disregard for the harm suffered by disordered gamblers, the shape of the Social Indifference Curves would be vertical, as illustrated in Figure 5. In this case, only R would matter, so any move rightward constitutes a better situation for society, regardless of the level of H. Of

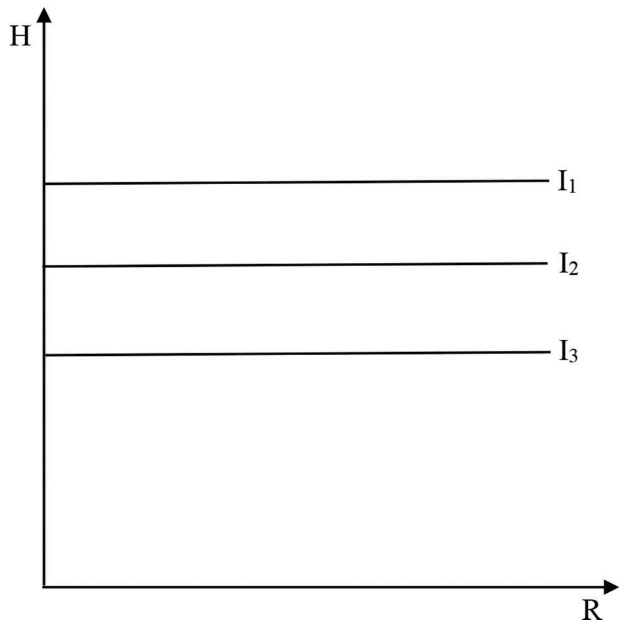

Figure 4. Three Social Indifference Curves that ascribe no social value to R; only the level of H matters to society.

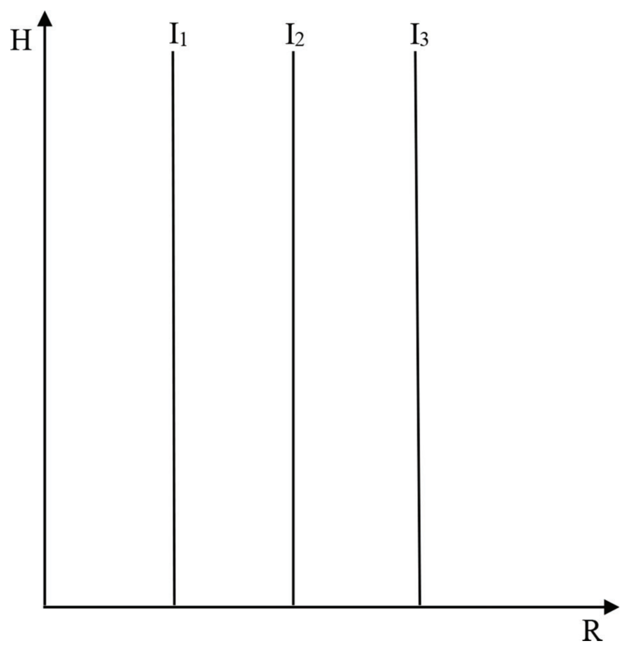

Figure 5. Three Social Indifference Curves that ascribe no social value to H; only the level of R matters to society.

course, the Social Indifference Curves having a shape that lies between these extreme cases are far more plausible and likely. The shape simply reflects society's values which, in a reasonably cohesive society, would express a social consensus.

Of course, Social Indifference Curves are a theoretical construct. When an analyst descends from the rarefied world of theory to the real world, it is admittedly impossible for a society to determine how its members rank every possible combination of the government's incremental net revenue and the harm inflicted on disordered gamblers. However, the point is that given the feasible options, *if* we could know the preferences of society at large, we could then determine the optimal feasible combination of H and R.

The Optimal Efficient Combination of H and R in the Long Run

What we have previously shown represents the optimal outcome for society in the "short run"—defined as a period of time that is insufficient to bring about any conceivable structural change in the gambling environment. The "long run" is defined—again functionally, not temporally—as a period of time that is sufficient to implement any desired change in the gambling environment (for example, altering the design of casinos or reconfiguring VLTs to provide cautionary pop-up messaging).[3] Thus, in the long run, by implementing various structural measures, the government could enlarge the set of feasible outcomes, thereby shifting the B_0 curve to the right to B_1, as in Figure 6 below.

In the long run, the set of feasible combinations of H and R has been enlarged to include all points on or above the B_1 line. The B_1 curve depicts the locus of all efficient combinations of H and R in the long run. The maximum attainable level of social welfare in the long run would occur along the B_1 curve at point E where the Social Indifference Curve I_3 is tangent to the B_1 curve. Comparing the optimal short run outcome at C with the long run outcome at E in Figure 6, society would achieve greater incremental net revenue as well as a reduction in the level of harm to disordered gamblers in the long run, as compared to the short run.

It is important to note that better results in the long run are not a foregone conclusion. If the government is incompetent or makes a misjudgment, it could end up making a long-run change that would move society to an inefficient combination of R and H. To illustrate this point, suppose the government decides to build a new casino in an area where the market for gambling products is already saturated. Because gambling consumption and H would remain constant, yet R would be reduced (the costs of the new useless casino would lower incremental net revenue), society would move from a point on the B_0 boundary to an inefficient point, even though there was the opportunity to move to the B_1 boundary. In this way, inefficiency could conceivably occur in both the long run and the short run. Just because more options are available does not necessarily mean the government will take full advantage of them for the betterment of society.

Applying the Framework to the Real World

Now that the conceptual framework for analyzing the trade-off between social harm and the government's incremental net revenue has been built, we can turn to the data

[3]This distinction between the short and long run stems from the work of the leading English economist of the early 20th century, Alfred Marshall.

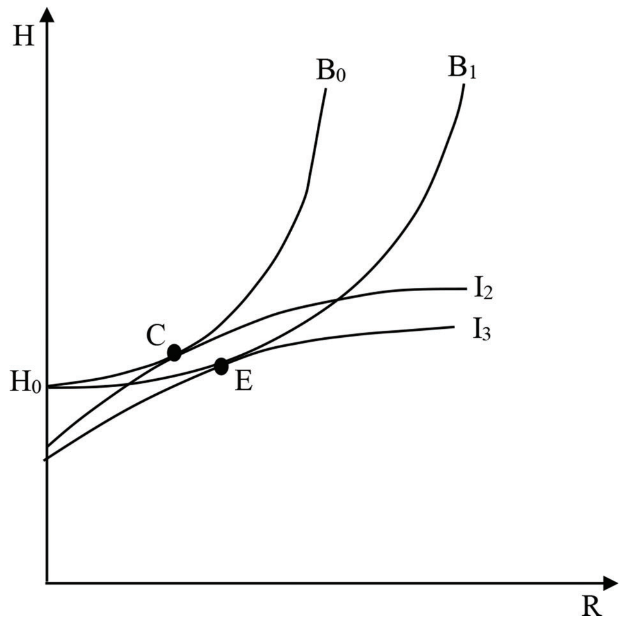

Figure 6. B_0 is the efficient boundary of R and H combinations in the short-run; B_1 is the efficient boundary of R and H combinations in the long-run. C is the optimal combination of R and H in the short-run and E is the optimal combination in the long-run.

to see how our model might appear in real life. Specifically, we will qualify the shape of the B_0 curve with evidence from the literature.

Although it is quite challenging to estimate the percentage of gambling revenue that is directly attributable to disordered [4] gamblers, evidence does exist to suggest that this group contributes disproportionately to casino revenues (Henriksson & Lipsey, 1999, p. 265). In Alberta, excessive spending on gaming by disordered gamblers totalled $326 million (ANIELSKI Management, 2008, p. 70).[5] Such expenditure amounts to an extra $2,890 in net expenditure per disordered gambler in the province (ANIELSKI, 2008, p. 69). When one considers that 3.2% of the adult population in Canada is ranked either as a moderate or severe "problem gambler" on the Canadian Problem Gambling Index (Wood & Williams, 2009, p. 7), this additional revenue becomes quite significant.[6] Given the disparity in spending between disordered and recreational gamblers and the relatively high prevalence of the former group in Canada, one could surmise that gambling revenues are quite dependent on the losses suffered by disordered gamblers.

[4]Whereas earlier studies have used the term "problem gambler," we have replaced this term by "disordered gambler," in keeping with the current terminology.

[5]"Excessive spending" is defined as "the difference between the net expenditures of problem gamblers and the average net expenditures of recreational (non-problem) gamblers" (ANIELSKI , 2008, p. 69).

[6]It should be noted that the prevalence rate of disordered gambling varies quite significantly from source to source. The most recent Canadian Community Health Survey that addressed disordered gambling estimated that the prevalence rate of problem gamblers and gamblers at a moderate risk of harm from gambling was 1.1% (Statistics Canada, 2014).

Of course, the harm suffered by disordered gamblers does not take the form of their monetary losses exclusively. Many would argue that the social costs of disordered gambling are at least as severe as the financial ones. Using a "Health Utility Index" that ranged from 0 to 1, research based on Swiss data concluded that disordered gamblers scored lower than the general population by 0.12, indicating a significantly lower health-related quality of life (Kohler, 2014, p. 10).[7] The same study also showed that disordered gamblers were more likely to be smokers, have a chronic disease or physical disability and were twice as likely to suffer from depression (Kohler, 2014, p. 10). This is in line with the data from Canada. According to the Canadian Community Health Survey microdata from 2013-2014, 17.6% of "problem gamblers" would also be considered depressed,[8] compared with only 5.3% of the rest of the gambling population (Statistics Canada, 2014). Although the total social cost to an individual disordered gambler is difficult to ascertain, estimates (from the 1990s) range from $13,200 to $56,000 per person if these qualitative costs were to be put into dollar values (Henriksson & Lipsey, 1999, p. 265). Total harm is therefore much more than only the money "transferred" to the government from the losses of disordered gamblers.

How much does the tax-paying public benefit from state-sponsored gambling? Using the government's incremental net revenue as our measure, "the ability of casinos to generate net benefits is marginal" (Henriksson, 1996, p. 117). Because most of the tax revenue collected from casinos could have been tax revenue collected from other businesses that have been pushed out by the casino, little new wealth is created for the government. New tourism generated from the opening of a casino could create wealth, but apart from the casinos in Windsor and Niagara Falls, and the gambling facilities in the Ottawa area,[9] there is no evidence to suggest that Canadian casinos have changed tourism profits at all (Henriksson, 1996, p. 117). Thus, the profits reported by casinos and the government can be misleading in estimating the actual benefit realized by society. This does not mean the government cannot realize positive incremental net revenue from operating casinos. However, any government revenue figures that ignore the tax revenue foregone through the crowding out of local businesses by casinos are significantly overstated.

Evidence from the literature seems to suggest an unsatisfactory trade-off between revenue and harm. Revenue is dependent on harm since disordered gamblers make up such a large percentage of gambling profits, with total harm being higher than simply the money lost to the government once social costs are taken into account. In

[7]The value of the Health Utility Index was determined by an algorithm that assessed problem gamblers' responses to two surveys, the SF-6D and the EQ-5D (Kohler, 2014, p. 6). It is beyond the scope of this paper to provide a detailed description of the construction of Kohler's Health Utility Index.

[8]This statistic is derived from the percentage of respondents that were both flagged as "problem gambler" by the survey and had a predicted probability of depression of 90%.

[9]Although there is no proper casino in Ottawa, there is the Rideau-Carleton Raceway Casino in Gloucester, as well as Casino du Lac-Leamy in Gatineau, Quebec.

addition, since we have defined R as incremental net revenue (having deducted tax revenue foregone from other sources with the advent of government-operated gambling), the magnitude of R may be relatively insignificant. The body of research therefore suggests that the B_0 curve in our conceptual framework should be quite steep, indicating that society must endure considerable H for a very modest increase in R. Such a trade-off between H and R would be reflected by the shape of the B_0 curve depicted in Figure 7.

There are certainly parts of the B_0 curve that remain contentious. For example, although the severity of the financial and social costs are quite clear from the literature, what is unclear is what portion of these costs are attributable to legalized gambling, and if social harm from gambling and the government's incremental net revenue necessarily move in the same direction on the efficiency boundary. Figure 8 illustrates a world where harm is actually reduced by the initial entry of government as the monopoly operator of gambling, thereby displacing (or driving out) gambling operations by organized crime. This would be reflected in the portion of the B_0 curve between H_0 and F. However, beyond point F, as the government engages in marketing to promote gambling activity to realize greater profit, social harm could eventually surpass the level it was at when all gambling activity was illegal. Notice that in Figure 8, in contrast to other diagrams in which the B_0 curve is depicted, the point at which social harm is minimized is not at the intercept H_0, but rather at F. While the entire B_0 curve in Figure 8 would delimit the boundary of feasible H and R combinations, the portion of the curve beginning at point F and to the right of F would represent the locus of efficient combinations of H and R. Figure 8 depicts a speculative B_0 curve, in the nature of a hypothesis that remains to be tested by empirical research in light of available evidence.

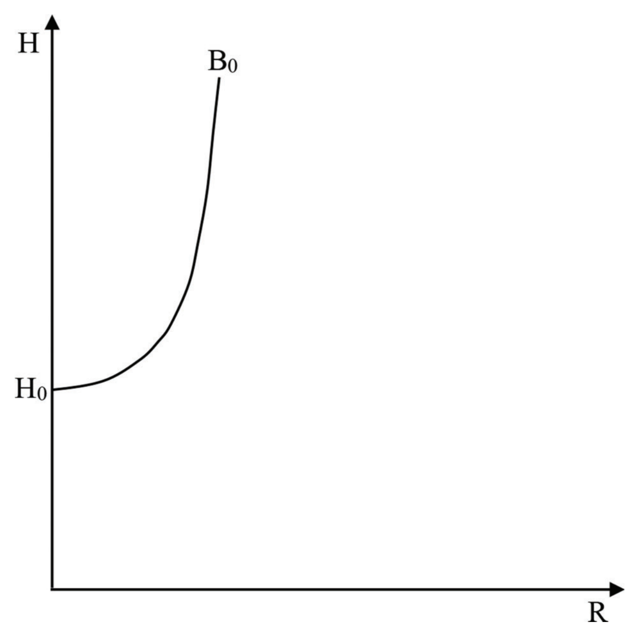

Figure 7. The efficient boundary B_0, where the steep slope accords with empirical evidence cited in the paper.

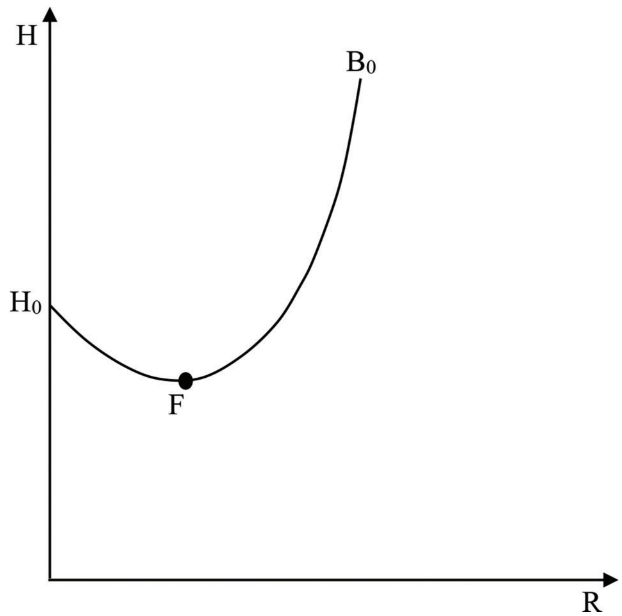

Figure 8. The efficient boundary B_0, where the H_0-F portion reflects a falling H as R increases. This accords with government-operated gambling displacing more harmful illegal private gambling.

Conclusion

To specify "harm minimization" as the goal of government-operated gambling is inappropriate. The goal of society should properly be stated as harm minimization for a given level of a government's incremental net revenue realized from gambling activity, or as maximizing the level of the government's incremental net revenue from gambling for a given level of social harm resulting from gambling activity. Either specification would imply "efficiency," which means being on the B_0 curve in the short run and the B_1 curve in the long run. We can determine this efficient B_0 curve by examining available data to gain an understanding of the general nature of the trade-off. There is evidence to suggest that the slope is quite steep, indicating a poor trade-off between H and R. While combinations of H and R on the B_0 curve are efficient, only one point on the B_0 curve is optimal, determined by society's preferences that are expressed by Social Indifference Curves. Whether government-operated gambling in its current form meets the standard of efficiency, let alone that of optimality, is a topic for debate. What is clear is that a conceptual framework is needed to uncover the trade-off between the social harm caused by government-operated gambling and the benefits derived from the incremental net revenue of that government.

References

ANIELSKI Management (2008). *The Socio-Economic Impact of Gambling (SEIG) Framework: An assessment framework for Canada: In search of the gold standard.* Np: np. Retrieved from: https://prism.ucalgary.ca/bitstream/handle/1880/45235/

seig_framework.pdf;jsessionid=A9B80E32D3B37F5256B62FE6269DC61D?sequence=1

Blaszczynski, A., Sharpe, L., & Walker, M. (2001). *The assessment of the impact of the reconfiguration on electronic gaming machines as harm minimisation strategies for problem gambling*. Sydney, Australia: University of Sydney Gambling Research Unit.

Dickerson, M. (2003). "Exploring the limits of 'responsible gambling': Harm minimisation or consumer protection." *Gambling Research, 15*, 28–43. Retrieved from: https://www.ipart.nsw.gov.au/files/a0cc0ae8-eaab-4767-9d46-9f26017e68ff/M_Dickerson_-_S5505.pdf

Henriksson, L. (1996). Hardly a quick fix: Casino gambling in Canada. *Canadian Public Policy, 22*, 116–128. doi:10.2307/3551903

Henriksson, L. & Lipsey, R. (1999). Should provinces expand gambling? *Canadian Public Policy, 25*, 259–274. doi:10.2307/3551892

Jackson, A. C., Christensen, D. R., Francis, K. L., & Dowling, N. A. (2016). Consumer Perspectives on Gambling Harm Minimisation measures in an Australian Jurisdiction. *Journal of Gambling Studies, 32*(2), 801-822. doi: 10.1007/s10899-015-9568-4

Kohler, D. (2014). A monetary valuation of the quality of life loss associated with pathological gambling: An application using a health utility index. *Journal of Gambling Issues, 29,* 1–23. doi:10.4309/jgi.2014.29.9. Retrieved from: http://jgi.camh.net/index.php/jgi/article/view/3904/4011

Mishan, E. J. (1980). How valid are economic evaluations of allocative changes? *Journal of Economic Issues, 14*, 143–161. Retrieved from http://www.jstor.org/journal/jeconiss

Quinn, F. L. (2001). First do no harm: What could be done by casinos to limit pathological gambling. *Managerial and Decision Economics, 22*, 133–142. doi:10.1002/mde.1003

Statistics Canada (2014). *Canadian community health survey: 2014–2015: Annual component*. Retrieved from http://odesi1.scholarsportal.info/webview/index.jsp?object=http://142.150.190.11:80%2Fobj%2FfStudy%2Fcchs-82M0013-E-2013-2014-Annual-component&mode=documentation&v=2&top=yes

Wood, R. T., & Williams, R. J. (2009). *Internet Gambling: Prevalence, Patterns, Problems and Policy Options.* Final Report prepared for the Ontario Problem Gambling Research Centre, Guelph, Ontario. January 5, 2009.

****

Submitted May 10, 2016; accepted September 16, 2018. This article was peer reviewed. All URLs were available at the time of submission.

For correspondence: Irwin Lipnowski, Ph.D., Department of Economics, University of Manitoba, 503 Fletcher Argue Building, Winnipeg, Manitoba, Canada, R3T 2N2. E-mail: Irwin.Lipnowski@umanitoba.ca

Competing interests: None declared (all authors).

Ethics approval: Not required.

Acknowledgements: The authors would like to thank Dr. Arman Vahedi for his help in the preparation of the diagrams, and Mr. Richard Lennon, who allowed us to draw upon his intimate knowledge of the gambling industry in Canada.