Economic Incentives for Protecting Digital Rights Online

N. Boris Margolin[a,*], Brian Neil Levine[b], James D. Miller[c], Matthew Wright[d,*]

[a] *73 Spring Park Avenue, Jamaica Plain, MA 02130, USA*
[b] *140 Governors Drive, Amherst, MA 01003, USA*
[c] *Pierce Hall 101, Smith College, Northampton, MA 01063, USA*
[d] *Department of Computer Science and Engineering, The University of Texas at Arlington, Arlington, TX 76019, USA*

Abstract

Once electronic content—such as a password to access a website's resources—has been released, it is very difficult to prevent it from being shared, which can cause economic harm to the content's owner and others. Most attempts to prevent unauthorized sharing of digital content have been based on technology or legal punishments, but these approaches have not been completely effective. We propose the use of economic incentives to both limit and detect unauthorized sharing. This approach has the advantage of not requiring watermarking, encryption, or other traditional Digital Rights Management techniques. Our protocol, called SPIES, is applicable to content that is shared to a limited extent and that meets several economic conditions. These conditions apply for many forms of content that are currently protected using technological sharing-prevention techniques. Such applications include passwords, trade secrets, pre-release content, and many others. We use formalize this protocol using game theoretic analysis, and we how to set the specific parameters under which SPIES can be useful.

Key words: digital rights management, security, incentives, game theory

*Corresponding author
Email addresses: `boris.margolin@gmail.com` (N. Boris Margolin), `brian@cs.umass.edu` (Brian Neil Levine), `EconomicProf@yahoo.com` (James D. Miller), `mwright@cse.uta.edu` (Matthew Wright)

1. Introduction

As websites offering content seek effective ways to generate revenue, a popular choice has been to provide access to its information only to subscribers. Sometimes the website protects most or all its content this way, as with Angie's List and Consumer Reports, while on other sites, subscribers pay for specialized information, as with ESPN. In both of these models, users are critical to the protection of the resources on the site, as they control an account name and password that can easily be shared, resold, or stolen from unsecured computers.

More generally, content with economic value is increasingly being offered in digital form and transferred electronically to collaborators, reviewers, distributors, and consumers. Once content has been distributed, it is very difficult for the owner to prevent unauthorized sharing. Content providers have tried to limit unauthorized sharing through the use of trusted software and hardware and other technological means, but these techniques have not been completely effective, as evidenced by the continuing availability of new copyrighted content on file-sharing networks.

We propose using *direct economic incentives* to discourage illegitimate sharing and to detect it when it occurs. In our scheme, users have a financial interest in keeping their content private. Direct economic incentives have the advantage of not being dependent on encryption, watermarking, or trusted hardware or software, while they can be used in conjunction with them.

An apparent disadvantage of economic incentives, when compared to technological restrictions, is that economic incentives only makes sharing less attractive and cannot entirely prevent it. However, as discussed by Haber et al. (2003), most existing technological content-protection measures have been circumvented. Thus, they also discourage sharing rather than prevent it. Moreover, computer security depends on layered defenses, and we do not intend our protocol to replace non-disclosure agreements, corporate policies and procedures, or other legal, technical, or physical protection layers. Rather, we seek to add economic incentives as an additional and complementary layer of protection.

Our protocol, called *Secret Protection Incentive-based Escrow System* (SPIES), is best suited for applications where the content is legitimately shared among a small set of persons. Participants first register with the content owner. Thereafter, anyone who has access to the content can use it to contact with a trusted server; the first such person receives a bounty payment. If the bounty is claimed, each legitimate content possessor is assessed a fine. Therefore, participants have an incentive not to share the content with anyone they do not trust. A major limit to the effectiveness of the protocol is that if a participant fears that someone else will share the content, despite the incentives against doing so, she may share the content defensively, trying to register first and sell the content to minimize her losses. We address this limitation by showing how to make create an environment in which cooperation in the protocol is more likely.

We formalize these intuitions using techniques from microeconomics. When the bounty and fine are set correctly, participants benefit most by keeping the content confidential. If for some reason a participant does decide to share the content, it is in her interest to first receive the bounty by registering the content, which has the positive side effect of signaling her intentions to others. One reason for sharing is fear that other participants may have shared the content themselves. This fear can be reduced, but not eliminated, by decreasing the benefits of registering and by increasing trust between participants.

For these properties to be achieved, our analysis shows that the bounty and fine must be set so that all participants, including the content provider, expect to be better off by participating in SPIES. Unfortunately, for some types of content, this is not always possible. If the participants can make a great deal of money by selling the content, it will be difficult to discourage them from doing so using fines. On the other hand, if the content is not particularly valuable to the participants, they will balk at risking a fine to access it. Similarly, if the content provider isn't much harmed by unauthorized sharing, she may as well just sell it. We formally define the exact conditions for the use of SPIES, which depend on the value of the content to the provider and the consumers, the salability of the content, and the harm to the provider of unauthorized sharing. The conditions are also modified by the degree of trust the participants in the protocol have in each other.

We also provide several variations of SPIES. For content that is valuable for a limited period of time, security deposits can be used instead of fines. Incentives for preemptive registration of content can be reduced using watermarking. In some cases, these incentives can be completely eliminated by keeping the content (like the user's password) hidden from the provider (who might instead have a hash of the password).

The remainder of the paper is structured as follows. In Section 2, we discuss related work in digital rights management, peer-to-peer systems, and economics. In Section 3, we present SPIES. We give example applications of SPIES in Section 4. In Section 5, we formalize the outcomes of SPIES and the strategies of the participants using game theory. In Section 6, we show how the content provider should set the fine and bounty, and under what economic conditions a suitable setting is possible. We present variations of SPIES in Section 7 and conclude in Section 8. The Appendix summarizes basic concepts in game theory.

Portions of this paper have appeared earlier as Margolin, Wright, and Levine (2004b,a).

2. Related Work

Preventing illegitimate sharing and using economic incentives in peer-to-peer protocols have both been popular areas of research in recent years.

Digital Rights Management (DRM) is the use of technological means to allow only authorized uses of or access to content. Traditional DRM uses techniques such as encryption (Sibert et al., 1995) and watermarking (Cox et al., 1996; Linnartz, 1998) in hardware or software to attempt to prevent acquisition or use of media unless the proper rights have been granted. SPIES, on the other hand, deals only with unauthorized sharing and uses explicitly financial means to discourage it.

As Haber et al. (2003) discuss, no DRM system to date has been able to completely prevent unauthorized access and use. Instead, DRM systems make it more inconvenient, expensive, or legally risky to access or use the protected content without authorization, creating implicit economic disincentives for the undesired behavior.

Horne et al. (2001) present the protocol most similar to SPIES. In their system, users are given incentives to redistribute content within a small peer-to-peer system, rather than incentives to not share it as in SPIES. If the content leaves the community, users stand to lose these payments.

Several protocols have been proposed to deal with the problem of insufficient sharing in peer-to-peer networks. This problem, known as *free-riding*, occurs when users try to get the benefit offered by a network (for example, quickly downloading files) without contributing their own work (serving files). Golle et al. (2001) construct several schemes to deal with the free-rider problem in peer-to-peer networks, thereby encouraging content sharing. In contrast, SPIES discourages content sharing.

We use the techniques of game theory to analyze strategies and outcomes in SPIES. Gibbons (1992) gives a good introduction to the subject. The strategies of participants in SPIES correspond to the strategies in the Assurance Game; the chief problem for such games is improving trust and coordination between the players. In the Appendix, we provide a summary of basic concepts of game theory, the Prisoner's Dilemma, and Assurance Games.

Farrell (1984) suggests that non-binding communication (*cheap talk*) can improve outcomes in some games. Arvan et al. (1999) show that when certain conditions are met, as in the Assurance game, cheap talk *must* improve outcomes. Baliga and Sjöström (2004) address the issue of achieving trust when the players are not certain about each other's rationality and preferences. Operating in the context of arms races, they again show that cheap talk can be used to convey information about one's outcome preferences, and that such information improves outcomes.

Our protocol employs a number of well-understood cryptographic techniques, including anonymous communication (Berthold et al., 2000; Freedman and Morris, 2002; Syverson et al., 1997; Reiter and Rubin, 1998), fair exchange (Bao et al., 1998), bit commitment (Schneier, 1996, pp. 87–88), and fair coin flips (Schneier, 1996, p. 89).

3. Protocol Description

In this section, we detail the operation of SPIES.

3.1. Variables in SPIES

SPIES is a protocol executed between several actors:

- A: Alice, the content provider
- B: Bob, the content consumer
- E: an escrow service
- U: Una, an unauthorized possessor of the content
- Z: a set of charities.

Alice and Bob do not trust each other to behave honestly, but both trust the escrow service to faithfully follow the protocol. Alice and Bob do not necessarily trust the escrow service with the protected content. Una and the charities do not trust each other or any other actor in the protocol.

The content is denoted ϕ. We use a *protection period* that lasts for time τ, which can be infinite. The fine is $\$f$, and the bounty is $\$b$. The semantic description and serial number of ϕ (e.g., *"password for account* `bmargolin` *for service to Consumer Reports, ID# 1234567890"*) is denoted d_ϕ.

See Figure 1 for a summary of the variables in SPIES.

In the protocol, $C \rightarrow D : \$x$ denotes an electronic payment of x dollars by C to D. The exchange of funds can be done by any secure method, such as electronic cash or credit cards over SSL. $(\$x)_C^D$ represents a commitment from C to pay $\$x$ to D if certain conditions are met. This commitment can take several forms, such as:

Variable	*Description*
ϕ	The protected content
A	The provider of ϕ
B	The consumer of ϕ
E	The trusted escrow service
U	An unauthorized possessor of ϕ
d_ϕ	A human-readable description of ϕ
$H(\phi)$	The secure hash of ϕ
τ	The time that protection of ϕ lasts
Z	The set of charities
$\$f$	The fine for sharing ϕ
$\$b$	The bounty for registering ϕ
$(\$x)_C^D$	A commitment of $\$x$ from C to D

Figure 1: Variables used in SPIES

- A physical, signed agreement establishing legal liability;
- An electronic document, signed with the participant's private key, establishing legal liability;
- An electronic cash payment, encrypted using the content (so that the cooperation of a content possessor is needed to decrypt it).

For the sake of clarity, we omit the fact that, where necessary, each party's message is signed for authenticity and integrity using previously exchanged public or shared keys.

3.2. Protocol Details

The SPIES protocol consists of an *Exchange* step followed by a *Protection Period*, which ends only if the Bounty is claimed.

1. During the **Exchange** period, the content consumer places a certain sum of money *at risk* by agreeing to pay that sum in the event that the protected content is shared. She then receives the protected content. To protect against dishonest content providers, the provider of the password places a sum of money at risk as well.

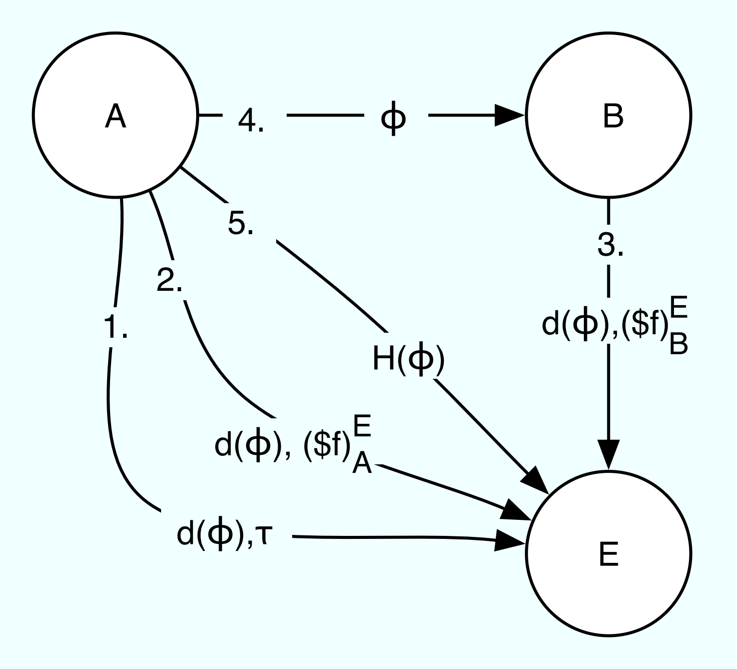

Figure 2: Steps 1–5: the *Exchange* period of the protocol.

2. During the **Protection Period**, anyone in the general public can register with an escrow service by providing proof of possession of the content. The first person to register is given a bounty payment. Anonymous registrations are accepted to encourage participation by those who fear prosecution.
   If the protection period ends and no one has claimed the bounty payment, the consumer's liability is erased. Otherwise, she must pay the fine. The bounty payment is strictly less than the fine, so the content consumer loses money even if she claims the bounty herself before sharing the content.
   Since the bounty is less than the fine, the escrow service has extra money if any registrations occur. This money is distributed to a randomly selected charity from a set agreed upon by the content provider and content consumer.

The protection period can be either limited or indefinite, depending on the characteristics of the protected content.

We now describe these stages in detail.

Exchange. These five steps are illustrated in Figure 2.

Step 1. $A \rightarrow E : d_\phi, \tau$

Alice registers the description of the secret d_ϕ and the protection period τ with the escrow agent E.

Step 2. $A \rightarrow E : d_\phi, (\$f)_A^E$

Step 3. $B \rightarrow E : d_\phi, (\$f)_B^E$

Bob and Alice place $\$f$ at risk by agreeing to pay this fine if sharing occurs.

Step 4. $A \rightarrow B : \phi$

Alice then sends Bob a copy of ϕ. If A or B requests it, Steps 2–4 can be done in a fair manner, such as with a protocol for fair exchange (Bao et al., 1998), so that Alice shares ϕ only when Bob has placed $\$f$ at risk, but we do not require a specific mechanism. A simpler approach would be to trust the escrow service to send Alice and Bob the other party's commitment to pay the fine.

Step 5. $A \rightarrow E : H(\phi)$

In the final step, Alice sends a secure hash of the content, denoted $H(\phi)$, to E. The use of a hash function in SPIES enables automated processing, but limits the protocol to contexts in which the content must be exact, such as a password. We discuss the use of SPIES in other contexts in Section 4.3. Alice's participation can be validated by having E send B the hash value. However, SPIES should only be used when Alice has an incentive to protect the content shared with Bob and therefore has no reason to cheat at the exchange step.

At the end of Step 5, Alice and Bob have $\$f$ each at risk to be paid if ϕ is revealed, and both Alice and Bob have knowledge of the secret ϕ (Figure 2).

Protection Period. Using an out-of-band means such as a public website, E publishes widely that it is seeking anonymous registrations from any unauthorized persons holding the protected content described by d_ϕ. If the protection

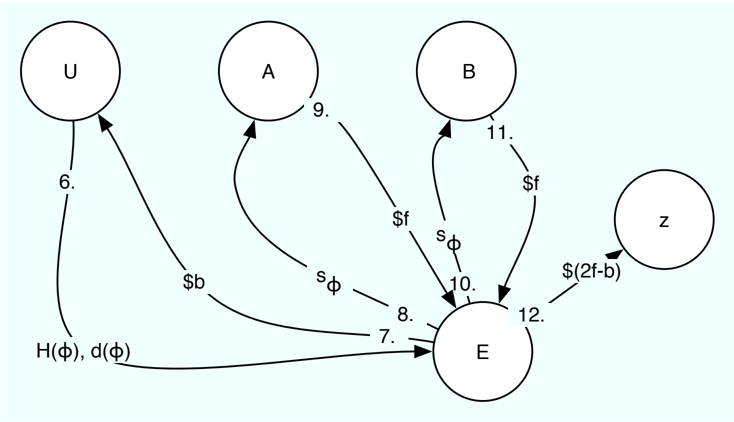

Figure 3: Steps 6–12: the *Bounty Payment*.

period is finite and no registrations are received before it ends after time τ, Alice and Bob's financial liabilities are erased. Otherwise, Steps 6–12 take place, as illustrated in Figure 3.

Step 6. $U \rightarrow E : H(\phi), d_\phi$

Una proves she has the content to the escrow service. As we show in Section 5.1.1, it is in the interest of any unauthorized possessor Una to register and try to get the bounty payment. Alice and Bob could register, but it is not in their interest to do so unless they intend to share the content.

Step 6. $E \rightarrow U : \$b$

Step 7. $E \rightarrow A : s_\phi$

Step 8. $A \rightarrow E : \$f$

Step 9. $E \rightarrow B : s_\phi$

Step 10. $B \rightarrow E : \$f$

Step 11. $E \rightarrow U : \$b$

If a correct registration is received, Alice and Bob are immediately notified and fined $\$f$ and the registrant is paid $\$b$. The message that indicates that ϕ has been shared is denoted s_ϕ. Steps 6 and 6 should be done in a fair manner (Bao et al., 1998) so that Una can have confidence that she will be paid.

Step 12. $E \rightarrow z \in Z : \$(2f - b)$

Finally the escrow service sends the remaining money to a randomly selected charity $z \in Z$. If desired, a secure random coin flip protocol (see examples in Schneier, 1996, p.89) between Alice, Bob, and the escrow service can be used to choose a charity. Otherwise, the escrow service simply selects a charity at random. Having a large set of charities makes it more difficult for a content possessor to collude with a charity. If there is only a single charity in Z, either Alice or Bob can make a deal to register the content and split the fines with the charity, defrauding the other player.

Alternatively, either can release the content as an indirect method of donating to his or her favorite charity. Increasing the set Z decreases the chances that an authorized content holder will collude with the charity chosen to receive the surplus fine money.

In addition to providing an incentive against unauthorized sharing, the fine serves as an incentive to improve security and prevent access to the secret. For example, a user may be trusted to not sell his website access on the black market, but may still need an incentive to keep the security on her computer at a high enough level to minimize the likelihood of stolen passwords.

SPIES does not provide any recompense to the content owner if content is released. Doing so would give an incentive to the owner herself to register with the escrow service.

4. Example Applications

SPIES is suitable for a variety of applications involving content that is legitimately available to a small number of people but should not be shared further without authorization. For SPIES to perform properly, certain economic requirements must be met: the content must have value to the participants, there must be limits on the profits available from unauthorized sharing, and the harm to the content provider when content is shared must not be too high. The exact economic requirements are discussed in detail in Section 6.

4.1. Passwords and Shared Keys

Many companies offer on-line subscription services providing news, financial information, real estate listings and the like. Passwords to these services are often shared, costing the companies potential customers, and it is difficult for this sharing to be detected. These passwords can be protected with the SPIES protocol. SPIES can also be used to protect any shared symmetric key for encrypted communication.

4.2. Corporate Collaboration

SPIES is well-suited to protecting the sharing of secrets between a few corporations, for instance to work on a joint project. Because corporations are likely to carefully consider the effects of their actions, they are less likely to defect, and this is known to all parties. In addition, they are likely to maintain a communications channel (discussed in Section 5.2), which further reduces the chance of defection. Naturally such sharing will be protected using legal mechanisms (e.g., non-disclosure agreements), but SPIES is a useful additional layer of protection. Further, if the bounty is set appropriately it can make quick detection of leaks much more likely.

4.3. Pre-release Media

SPIES can be used to protect movies, songs, photos, and articles that are released to a limited number of persons for review before being released to the general public. Since the content only needs to be protected for a limited period of time, security deposits can be used as an alternative to fines (see Section 7.1).

When considering protection for such media, the hash function is no longer appropriate as a means to check for copying, as the media can be slightly modified or entirely recoded (e.g. to a different file format) without losing the value of the content. We suggest several possible mechanisms to verify the unauthorized possession of the content.

- **Trusting the Escrow Service with the Content.** If all parties trust the escrow service with the content, then Alice can give the escrow service the content and any unauthorized possessor can provide a copy of the content that can be checked against the original. The amount of content to check should not be prohibitively large and can therefore be checked by people. Non-matching copies, especially, can be filtered out quickly.

- **Partial Sharing and Challenges.** If sharing all of the content with the escrow service is too risky, small pieces of the content could instead be shared. For example, the content provider could give the escrow service 100 random 10 second clips from a movie. The escrow service could then challenge any unauthorized possessor to produce a randomly selected subset of these clips, avoiding synchronization issues by sharing single frames from the middle of each clip. It is likely that the escrow service would have to distribute special software to facilitate this process.

- **Robust Watermarks.** Although watermarks may not be sufficient as a DRM measure, they can be combined with SPIES to create an incentive to not share content. A robust watermark including a specific random string can be embedded in the content. The escrow service is given the watermark string and a key used to find the watermark. The unauthorized possessor provides a copy of the content, and the escrow service determines if the watermark is present to see if it extracts the same string.

Robust watermarks are especially useful if multiple parties receive copies of the media for review. This helps to clearly distinguish between them. Multiple users might try to defeat the watermarking by combining multiple legitimately obtained copies with different watermarks. Using collusion-resistant techniques (Boneh and Shaw, 1995; Dittmann et al., 2000), Alice can still punish at least some of the users involved.

4.4. An Anti-Example: Widely Sold Media

To clarify when SPIES can be used, it is helpful to include an example where SPIES would not be suitable. When media is sold widely to a large number of users, such as music and movies at the iTunes store, SPIES is no longer viable. To make it plausible at all, each copy must be robustly watermarked with a long, random string. A break in the watermarking scheme would be very expensive for the content owner in possible fines. Even then, the benefit of using the scheme recedes quickly with a large number of buyers. One irrational buyer who sells the content or one insecure computer from which the content could be stolen means that the content could be widely distributed on the Internet. Further selling or distribution of the content after this would be much less important to the content owner's losses.

From this example, we can see that SPIES requires a small set of players to be successful (although more than two can participate – see Section 7.2). More sharing increases the risk of wide distribution, thereby reducing the benefit of SPIES to the content owner. SPIES also benefits from (though it does not require) cheap talk, which is easier among a small group than an impersonal setting. Finally, the timeliness of the content is often important. Passwords protecting access to a website are much more important if the content on the website is updated frequently (such as ESPN expert tips for fantasy sports leagues) or contains a large variety of information (such as Consumer Reports); such information is more difficult to reproduce in a useful way than fixed content, like a song or a movie. Similarly, pre-release media distributed to a small group of reviewers is much more timely than media currently available for sale online.

5. Game-theoretic analysis

SPIES is based on economic incentives, so it relies for its operation on participants' regard for their own best economic interests. In this section we demonstrate formally what is in the best interest of the various players in SPIES. We show that:

- Rational players prefer not to share the content.
- Rational players may share the content if they fear another player will do so.
- If a rational player shares the content (for whatever reason), she will first register with the escrow service, so sharing will be detected.

These properties depend on the way bounties are paid, and they do not necessarily hold if the bounty payments are shared with the first several registrants, or if they depend on the number of registrations received. Game theory is the study of processes whose outcomes depend on the actions of several independent actors. The actors prefer some outcomes to others, and they are selfishly rational: they attempt to make the choices that will lead to the best possible outcomes for themselves.

5.1. The SPIES Game

SPIES is a static game of complete information. This means that each player knows the other players' utility functions, not just their own, that players choose their strategies simultaneously, without knowledge of others' choices, and that the game is not repeated.

SPIES has three players, Alice, Bob, and an unauthorized player Una. However, as we show, Una prefers the same strategy no matter what Alice and Bob may do, so her choice can be assumed. When Una is not considered, SPIES is reduced to an Assurance game (see Appendix A.2) between Alice and Bob. The outcomes of games of this type depend on the confidence the players in their estimates of each other's rationality and outcome preferences. Likewise, in SPIES, a positive outcome depends both on the setting of the fine and bounty payment (which affect the players' outcome preferences) and on the level of confidence that can be established between the players.

In the following examples, we have set the bounty payment, the fine, and the value of the content to the players to appropriate numerical values. This is a common technique in game theory, since it greatly simplifies the analysis.

In Section 6, we do a full analysis, showing what the relationships between these variables must be. Readers, even those familiar with the concepts of game theory, may find it useful to briefly read Appendices A.1 and A.2 to see simple examples of the tables and terminology used below.

5.1.1. $4 \times 4 \times 4$ *Game*

The three participants in SPIES are Alice, Bob, and Una. There may be more than unauthorized possessor, but this does not change the analysis. Each of these participants has four strategies to choose from:

1. C: To cooperate and not share the content
2. R: To register with the escrow service without sharing the content
3. S: To share the content without registering with the escrow service
4. SR: To both share the content and register with the escrow service

For this example we assume a fine of \$6; a sales price of \$3; and a bounty of \$2. These values give the correct incentives in SPIES; we show in Section 6 what values are possible. In this case, it is key that the sales price plus the bounty is less than the fine: a player loses \$1 if she registers and sells the content, but then has to pay the fine. (Any extra harm to Alice when the content is shared is excluded for simplicity; it does not affect the analysis.)

We assume that Alice and Bob start the game with a utility of \$6, and Una starts with a utility of \$0. This would occur, for example, if Bob values the content at \$12 and pays Alice \$6 for it. We can assume that Alice and Bob have a positive utility from exchanging the content, since they would not participate in SPIES at all otherwise. In any case, starting utility does not alter the strategies of the players; it simply makes the analysis clearer by keeping utility values positive.

The following are the outcomes for the players, given the strategies chosen.

1. Selling by playing the strategies S or SR adds \$3 to the utility, assuming the player has access to the content. Alice and Bob have always access to the content; Una has access to it only if Alice or Bob sells (plays S or SR).

C	C	R	S	SR
C	(6,6,0)	(0,2,0)	(6,9,0)	(0,5,0)
R	(2,0,0)	(1,1,0)	(2,3,0)	(1,4,0)
S	(9,6,0)	(3,2,0)	(9,9,0)	(3,5,0)
SR	(5,0,0)	(4,1,0)	(5,3,0)	(4,4,0)

R	C	R	S	SR
C	**(6,6,0)**	(0,2,0)	(0,3,2)	(0,5,0)
R	(2,0,0)	(1,1,0)	(2,3,0)	(1,4,0)
S	(3,0,2)	(3,2,0)	(3,3,2)	(3,5,0)
SR	(5,0,0)	(4,1,0)	(5,3,0)	(4,4,0)

S	C	R	S	SR
C	(6,6,0)	(0,2,0)	(6,9,3)	(0,5,3)
R	(2,0,0)	(1,1,0)	(2,3,3)	(1,4,3)
S	(9,6,3)	(3,2,3)	(9,9,3)	(3,5,3)
SR	(5,0,3)	(4,1,3)	(5,3,3)	**(4,4,3)**

SR	C	R	S	SR
C	**(6,6,0)**	(0,2,0)	(0,3,5)	(0,5,3)
R	(2,0,0)	(1,1,0)	(2,3,3)	(1,4,3)
S	(3,0,5)	(3,2,3)	(3,3,5)	(3,5,3)
SR	(5,0,3)	(4,1,3)	(5,3,3)	**(4,4,3)**

Figure 4: $4 \times 4 \times 4$ SPIES Game. Succeeding matrices indicate Una's choices; bold entries indicate Nash equilibria.

2. If a player who has content registers, playing R or SR, he or she has a chance of getting the bounty payment. The bounty is \$2, so if only one player registers their own utility is increased by \$2. If both Alice and Bob register, their utilities are each increased by \$1: each of them is equally likely to have registered first, so their expected return from registering is half of the bounty. However, Alice or Bob can always register before releasing the content. Thus, Una gets no benefit from registering if Alice or Bob has registered. She also gets no benefit if neither has shared the content, since she cannot present the content to the escrow service.
3. If R is played by any player who has the content, Alice and Bob have to pay the fine, so \$6 is subtracted from both of their utilities.

These modifications are cumulative, so if Alice plays SR her utility is increased \$3 by Rule 1, increased by \$2 by Rule 2, and decreased by \$6 by Rule 3, for an overall decrease of \$1.

If we let X_D indicate that player D has chosen strategy X, and let u_D be D's utility, then we can describe these conditions formally as follows:

$$\begin{array}{llll}
u_A & = 6 & & \\
 & +3 \text{ } \mathit{if} & SR_A \vee S_A & \text{Rule 1} \\
 & +2 \text{ } \mathit{if} & (SR_A \vee R_A) \wedge \neg(SR_B \vee R_B) & \text{Rule 2} \\
 & +1 \text{ } \mathit{if} & (SR_A \vee R_A) \wedge (SR_B \vee R_B) & \text{Rule 2} \\
 & -6 \text{ } \mathit{if} & SR_A \vee R_A \vee SR_B \vee R_B & \\
 & & \vee((SR_U \vee R_U) & \\
 & & \wedge(SR_A \vee S_A \vee SR_B \vee S_B)) & \text{Rule 3} \\
u_B & = 6 & & \\
 & +3 \text{ } \mathit{if} & SR_B \vee S_B & \text{Rule 1} \\
 & +2 \text{ } \mathit{if} & (SR_B \vee R_B) \wedge \neg(SR_A \vee R_A) & \text{Rule 2} \\
 & +1 \text{ } \mathit{if} & (SR_B \vee R_B) \wedge (SR_A \vee R_A) & \text{Rule 2} \\
 & -6 \text{ } \mathit{if} & SR_B \vee R_B \vee SR_A \vee R_A & \\
 & & \vee(SR_U \vee R_U) & \\
 & & \wedge(SR_B \vee S_B \vee SR_A \vee S_A) & \text{Rule 3} \\
u_U & = 0 & & \\
 & +3 \text{ } \mathit{if} & (SR_U \vee S_U) & \\
 & & \wedge(SR_A \vee S_A \vee SR_B \vee S_B) & \text{Rule 1} \\
 & +2 \text{ } \mathit{if} & (SR_U \vee R_U) \wedge (S_A \vee S_B) & \\
 & & \wedge\neg(SR_A \vee R_A \vee SR_B \vee R_B) & \text{Rule 2}
\end{array}$$

We show the utilities for the various strategies of Alice, Bob, and Una in Figure 4.

The game has four Nash equilibria, which appear in bold: one at (C,C,R), one at (C,C,SR), one at (SR,SR,S), and one at (SR,SR,SR). There is no real difference between (C,C,R) and (C,C,SR), since Una's strategies have no effect unless she has access to the content.

There is no equilibrium for which Alice or Bob plays R, since SR dominates R for both of them. No matter what, Alice or Bob can get more by both selling the content and registering than by registering only, which is not surprising.

Una does not have any strictly dominated strategies. However, SR is never a worse choice for Una than C, S, or R, and, depending on Alice and Bob's

	C	*S*	*R*	*SR*
C	**(6,6)**	(0,2)	(0, 3)	(0, 5)
S	(2,0)	(1, 1)	(2,3)	(1, 4)
R	(3,0)	(3, 2)	(3, 3)	(3, 5)
SR	(5,0)	(4, 1)	(5, 3)	**(4,4)**

Figure 5: 4×4 SPIES Game

strategies, it may be better, so SR is a weakly dominating strategy. If Una has access to the content, she has at least a chance of being better off registering it, so it is in her interest to register. S appears as a strategy for Una in one of the equilibria since Una may assume that the content must have been registered.

There are no Nash equilibria in which content is shared but not registered by anyone, so given the fine, bounty, and sales price in the game, it is likely that sharing of content will be detected by the Escrow service.

5.1.2. 4×4 *and* 2×2 *Games*

Una can never go wrong by playing SR, so we assume that she will. Una's actions and utilities are implicitly those of the bottom matrix of Figure 4, representing her play of strategy SR.

We show the reduced SPIES game in Figure 5. Again we assume a fine of \$6, a sales price of \$3, a bonus of \$2, and initial utilities of \$6 for Alice and Bob. Again, the key is that the fine exceeds the sum of the bonus and the sales price.

The equations for Alice and Bob's utilities are the same as above. Since Una always plays SR, SR_U is always true.

For both Alice and Bob, SR strictly dominates both S and R. It is better for each of them to play SR than either S or R, no matter what choice the other player makes. In other words, if a player plans to sell the content, she should always register first with the escrow service. Registration is anonymous, so she will not be punished for doing so, although she will have to pay the fee, along with the other authorized possessor.

Since rational players never play strictly dominated strategies, we further simplify the reduced game to eliminate the strategies R and S for Alice and

	C	SR
C	$(\mathbf{6},\mathbf{6})$	$(0,5)$
SR	$(5,0)$	$(\mathbf{4},\mathbf{4})$

Figure 6: 2 × 2 SPIES Game

Bob. We show the result in Figure 6. The game has two Nash equilibria, one at (C,C) and one at (SR,SR). Of these, (C,C) is preferred since it leads to the highest utility for the two players. This graph has the same structure as the Assurance game, with C corresponding to Cooperate and SR corresponding to Defect.

Like other Assurance games the reduced version of SPIES is subject to the problem of distrust, so (SR,SR) is a likely outcome even with rational players, though both prefer (C,C). (SR, SR) is a less optimal result, but it is also less risky. By choosing SR Alice and Bob can ensure utility of at least $4, while choosing C risks ending up with a utility of $0.

5.2. Improving Trust

Since, from the point of view of Alice and Bob, the SPIES game corresponds to the Assurance Game (see Appendix A.2), some of the techniques that have been shown to improve outcomes in the Assurance game can be used to improve outcomes in the SPIES protocol.

There are three methods to improve outcomes in the Assurance game:

1. Change the game so that the parties can commit to Cooperating.
2. Reduce the penalty for Cooperating when the opponent Defects.
3. Improve trust between the participants. In this context, trust involves beliefs about each other's rationality and preferences, not their honesty.

Method 1 would be satisfied if, after the game starts, at least one of the players is actually unable to share the content, and both players know this. One way to accomplish this is through traditional, preventative DRM. Another way of preventing sharing is if only one of the players (Bob) actually has access to the content. This can be accomplished using cryptographic techniques such as Oblivious Transfer (Schneier, 1996, pp. 116-117) and Secure Multi-Party Computation (Chaum et al., 1988) to produce, for example, a password that only Bob learns or watermarked media that Alice can recognize, but not reproduce. Note that only the watermark need be private.

Method 2 can easily be applied to SPIES. The fine should be made large relative to the sum of the bounty and the sales price. Then Alice and Bob have little to gain, and much to lose, by sharing the content. Preventative DRM applies here as well; if, for at least one of the players, it is very difficult and expensive to share the content, the probability of the cooperative equilibrium is improved.

Method 3 consists of improving the trust of the participants in each other's rationality and shared preferences. Coordination between the players in an Assurance game can reduce the chance of the inferior equilibrium being selected. Even communication that does not commit a player to any particular course of action (cheap talk) can increase the likelihood of cooperation in small groups. Farrell (1984) and Arvan et al. (1999) show that cheap talk in some cases can, and in some cases must, improve the expected utilities of the participants. This communication improves players' knowledge of each other's preferences and trust in each other's rationality. For example, Baliga and Sjöström (2004) assume that there are two types of countries, Doves who prefer peace and use weapons to discourage attack and more aggressive but less common Hawks; cheap talk in arms negotiations can convey information about the the type of a country and improve all player's expected outcomes. Similarly, there may be two types of players in SPIES, Law-abiders and Sharers, who have different utility functions: Sharers gain some significant utility by sharing or selling the content, while Law-abiders do not. Communication during the protection period can increase confidence in player's assessments of each other's type.

For example, suppose Alice and Bob are companies sharing technical content so that they can collaborate. While the companies keep communicating, they are likely to believe that the other still wants the collaboration to success, and so has an interest in keeping the content confidential. In this case neither side will be likely to share the content merely out of fear that the other is planning to do so. On the other hand, if one side cuts off communication, the other may be afraid that the collaboration is doomed and will be more likely to register the content preemptively.

Publicly published information can constitute cheap talk. If a company has an elaborate web site discussing its content licensing business, a customer is likely to conclude that it is unlikely to share the content in order to get a bounty payment. Believing this, she is not likely to share the content preemptively. For example, a consumer can conclude that a movie studio making money by distributing digital versions of a popular film is unlikely to also post the film to a file-sharing network.

5.3. Bounty Payment Schemes

The bounty payments in SPIES encourage detection of unauthorized sharing, triggering a fine to the authorized possessors. These properties derive not just from the payment of a bounty, but from the way that it is paid.

To demonstrate the superior properties of the First Registration bounty payment scheme used in SPIES, we compare three possible bounty payment schemes:

- Exponential, as used in previous versions of SPIES (Margolin et al., 2004b,a): Bounty payments to individuals decrease exponentially with additional registrants.

- Limited: Bounty payments decrease with additional registrants, but only the first k registrations are considered.

- First Registrant, the scheme used in this paper: The first registrant with the Escrow service receives a bounty. Alice and Bob do not need register with the Escrow service.

As we show below, the advantage of the First Registrant scheme is that participants register if and only if they are going to sell the content to Una.

This acts as a signal to the content owner that the content is in the hands of unauthorized possessors.

5.3.1. Exponential

In the exponential bounty payment scheme, the bounty payment for each successive registration decreases exponentially. If there are two legitimate possessors of the content who each pay a security deposit of $\$f$ and n registrations (none are expected unless illegitimate sharing occurs), each registrant receives a bounty of

$$\frac{\$f}{2^n}.$$

If there are no registrations, Alice and Bob each receive their security deposits of $\$f$ back.

The exponential scheme can prevent any advantage to the Sybil attack (Douceur, 2002) of making multiple registrations using false identities. For each player, making k registrations strictly dominates making $k+1$ registrations, regardless of the actions of others. For legitimate possessors, the dominant strategy is making no registrations; for unauthorized possessors, the dominant strategy is making one registration.

However, this scheme has two problems. First, the use of a security deposit, rather than a fine, requires a limited protection period. Second, there is no difference in Alice and Bob's behavior vis-a-vis the Escrow service whether they intend to share the content or not, so unauthorized sharing is not detected right away. In contrast, in the first-registration bounty payment scheme it is in Alice and Bob's interest to register if, and only if, they are going to share the content; i.e., the strategy Sell & Register strictly dominates Sell. If they should sell without registering, it is likely that the buyer will immediately register the content. Thus it is in Alice and Bob's interest to register when they are about to sell, even though this signals unauthorized sharing.

5.3.2. Limited

A second scheme is to reduce bounty payments when there are multiple registrations, but to limit payments to the first few registrations. The disadvantages of this scheme are that it requires a limited protection period, and that registrations do not necessarily indicate that content has been shared.

For example, for a security deposit of \$8, the bounty for one registration might be \$7; for two registrations, \$3 each; and for three registrations \$1 each. For more than three registrations only the first three registrants would receive a payment of \$1. What is important is that the payment for $k+1$ registrations total is always less than the payment for k registrations, to discourage multiple registrations from the same person.

This scheme does discourage sharing of the content. However, making fewer registrations no longer strictly dominates and there are many more Nash equilibria: there are a number of cases where it is advantageous for a player to make multiple registrations in order to mitigate his or her harm from others' registrations.

Using the example security deposits and bounties given above and assuming that any sales profits are not significant relative to the fines and bounties, we found the following Nash equilibria for Alice and Bob:

- (No registrations, No registrations)
- (1 registration, 1 registration)
- (3 registrations, 3 registrations)
- (3 registrations and a sale, 3 registrations and a sale)

In contrast to the Exponential scheme, in this scheme there are equilibria in which a player registers even though he or she does not intent to share the content. Therefore registrations no longer indicate that the content has been shared.

Since bounties are not paid until the Protection Period is complete or three registrations have arrived, the Protection Period must be limited with this scheme.

5.3.3. First Respondent

In the First Respondent bounty payment scheme, if there are any registrations, the first registrant receives the bounty and the others receive nothing; Alice and Bob must pay a fine if a registration occurs.

This scheme has several advantages over the other two. First, it works whether the protection period is finite or not. Second, in this scheme, it is never advantageous for authorized possessors to make a registration when they do not intend to sell. When an authorized possessor does intend to sell, he or she has an incentive to signal this choice by registering immediately. Therefore registrations indicate likely sharing of the content.

The game analysis extends easily to multiple unauthorized possessors; since Alice and Bob are likely to register before selling, whether any unauthorized possessors register or not becomes irrelevant. If, for simplicity, we allow only the choices of Cooperating (C) or Selling & Registering (SR), there are two Nash equilibria at $(C, C, C, \ldots, C)$ and $(SR, SR, SR, \ldots SR)$.

6. Economic Requirements

In the previous section, we showed that SPIES prevents and detects sharing using specific numerical values to represent economic outcomes. We now determine what factors make up the utility of each outcome and the required economic conditions for SPIES to be used successfully.

SPIES protects content from unauthorized sharing by imposing a cost on sharers. Specifically, the bounty payment must be enough that it is worth the trouble for unauthorized possessors to register, and the fine must be enough to discourage unauthorized sales of the content. In some situations SPIES is effective, but superfluous. If a content provider isn't harmed by unauthorized content sharing, there's no incentive for her to use SPIES (or any other protocol that limits sharing). In addition, neither the content provider nor content consumers are likely to participate if they think that they will have to pay the fine, resulting in a net loss.

In this section we show under what economic conditions SPIES is applicable, and how the content provider should set the fine and bounty when it is.

There are three economic conditions on the use of SPIES:

Condition 1. *Both Alice and Bob must have a higher expected utility from playing C (keeping the content private) than from playing SR (sharing the content).*

Condition 2. *The expected utility of using SPIES must be positive for both Alice and Bob.*

Condition 3. *SPIES must result in higher expected utility for Alice that what she would obtain by selling the content without using SPIES.*

These conditions can be met when the content is valuable to Alice and Bob, when only a limited amount of money can be made by sharing the content without authorization, and when there is serious, but not fatal, harm to Alice from unauthorized sharing.

Variable	*Description*
ϕ	The protected content
f	The fine for sharing the content
b	The bounty payment
i_A, i_B	Alice and Bob's initial utilities
h_A	The harm to Alice from sharing
s	Profit from unauthorized sales
p_A	Alice's prob. of cooperating, as seen by Bob
p_B	Bob's prob. of cooperating, as seen by Alice

Figure 7: Economic Variables in SPIES

	C	SR
C	(i_A, i_B)	$(i_A - f - h_A, i_B + b + s - f)$
SR	$(i_A + s + b - f - h_A, i_B - f)$	$(i_A + s + \frac{b}{2} - f - h_A, i_B + s + \frac{b}{2} - f)$

Figure 8: General 2×2 game

The three conditions for SPIES to be effectively used hold when

$$\max\left(\frac{s}{p_B},\frac{s}{p_A}\right) < f < \min\left(\frac{i_A}{1-p_B} - h_A, \frac{i_B}{1-p_A}, \frac{h_A p_B}{1-p_B}\right) \tag{1}$$

and

$$b < \min\left(\frac{2p_B f - 2s}{1+p_B}, \frac{2p_A f - 2s}{1+p_A}\right) \tag{2}$$

are satisfied.

Alice is responsible for setting the fine and the bounty so that these inequalities are satisfied. This is possible only when

$$\max(\frac{s}{p_B},\frac{s}{p_A}) < \min(\frac{i_A}{1-p_B} - h_A, \frac{i_B}{1-p_A}, \frac{h_A p_B}{1-p_B}). \tag{3}$$

In the remainder of this section we derive these inequalities.

6.1. Expected Return

To decide both whether to participate in SPIES and what strategies to use, Alice and Bob must estimate, in terms of their net utility, what the cost or benefit will be.

The utilities of the 2×2 game of SPIES are shown symbolically in Figure 8. In addition to the bounty payment b and the fine f we consider Alice's initial utility, typically the sales price charged by Alice for the content, and denote it i_A. Bob's initial utility, typically his valuation of the content above Alice's sales price, is denoted i_B. We denote the harm to Alice when content is widely shared as h_A. We denote the amount available to Alice or Bob by selling the content without authorization by s. Note that the gain from selling must include such possibilities as sharing the content between a group of users, who in turn share other content with the sharing party, or selling the content to many other users. Nevertheless, Alice is likely to be able to make more than Bob by selling the content, since she is the content owner. We conservatively assume that she will make the same as Bob. Finally, we denote Alice's probability of cooperating, as estimated by Bob, as p_A, and Bob's probability of cooperating, estimated by Alice, as p_B. The variables in a content protection problem are summarized in Figure 7.

The utility that Alice and Bob expect from playing the SPIES game also depends on what probabilities they assign to the outcomes in the game, which depends on their beliefs and certainty in their beliefs about each other's preferences. These are very difficult to predict, even when all participants are rational and have the expected preferences.

Consider the Rock-Paper-Scissors game discussed in Appendix A.1. After Player 1 has played Rock several times in a row, should she play it again? She can reason that her opponent will expect this, and so play Paper. Therefore, perhaps she should play Scissors, which beats Paper. However, her opponent may anticipate this, and play Rock herself in anticipation to beat Scissors. Player 1 can anticipate this as well, and so on.

To avoid such recursive analyses, we assume that Alice and Bob, having thought the matter through, have arrived at probabilities p_A and p_B, where p_A gives Bob's estimate of Alice's probability of cooperating, and p_B gives Alice's estimate of Bob's probability of cooperating, and that these estimates already include all recursive considerations of each other's estimates.

If Bob believes that Alice has probability of p_A of playing C and probability of $1 - p_A$ of playing SR, then Bob's expected utility for playing C is given by

$$u_B(C) = i_B - (1 - p_A)f, \tag{4}$$

and Bob's expected utility for playing SR is given by

$$\begin{aligned} u_B(SR) &= p_A(i_B + b + s - f) + (1 - p_A)\left(i_B + s + \frac{b}{2} - f\right) \quad (5) \\ &= i_B + s - f + \left(\frac{p_A}{2} + \frac{1}{2}\right)b \quad (6) \end{aligned}$$

Likewise, if Alice believes Bob has a probability of p_B of playing C and a probability of $1 - p_B$ of playing SR, Alice's expected utilities for playing C and SR are given respectively by:

$$u_A(C) = i_A - (1 - p_B)(f + h_A) \tag{7}$$

and

$$u_A(SR) = i_A + s - f - h_A + \left(\frac{p_B}{2} + \frac{1}{2}\right)b. \tag{8}$$

Una's best choice is to always plays SR. Her expected utility is given by:

$$u_U(SR) = p_S(s + (1 - p_R)b) \tag{9}$$

where p_S is her estimate of the probability that someone will share the content and p_R is her estimate of the probability that someone will register the content.

6.2. Condition 1

The incentives provided by SPIES must make both Alice and Bob prefer to cooperate rather than to share the content.

Condition 1 is represented by the following inequalities. As above, we denote a player D's utility as a result of playing X by $u_D(X)$.

$$u_A(C) > u_A(SR) \tag{10}$$

$$u_B(C) > u_B(SR) \tag{11}$$

Both inequalities must be satisfied for the condition to be satisfied.

Inequalities (10) and (11) can be rewritten in terms of the expected utilities of the SPIES game, shown in Figure 8:

$$i_A - (1 - p_B)f - (1 - p_B)h_A > i_A + s + \left(\frac{1 + p_B}{2}\right) b - f - (1 - p_B)h_A \tag{12}$$

$$i_B - (1 - p_A)f > i_B + s + \left(\frac{1 + p_A}{2}\right) b - f \tag{13}$$

By solving these inequalities for the bounty b, we obtain:

$$b < \min\left(\frac{2p_Bf - 2s}{1 + p_B}, \frac{2p_Af - 2s}{1 + p_A}\right). \tag{14}$$

Alice must set the fine and bounty payment so that they satisfy this inequality. Bounty payments can be set arbitrarily low, but must be positive. Therefore the right hand side of (14) must be positive as well, which is true when:

$$f > \max\left(\frac{s}{p_B}, \frac{s}{p_A}\right). \tag{15}$$

Inequality 15 reflects that the fine must be set higher than any profit obtainable from unauthorized sharing for SPIES to be effective. The fine is also dependent on a player's estimate of the other player's probability of cooperating.

6.3. Condition 2

Both Alice and Bob must expect positive net utility from participating in SPIES or they are not likely to participate.

We now consider under what conditions SPIES gives a positive utility for Alice and Bob. We assume for simplicity that b is small enough to be ignored by Alice and Bob; it can be made very small, and it motivates Una to register as long as it is greater than zero.

For Condition 2 to hold, both

$$u_A(C) > 0 \tag{16}$$

and

$$u_B(C) > 0 \tag{17}$$

must be satisfied. Substituting the utilities from Figure 8, these become

$$i_A - (1 - p_B)(f + h_A) > 0 \tag{18}$$

and

$$i_B - (1 - p_A)f > 0. \tag{19}$$

Solving for f, these conditions reduce to

$$f < \min\left(\frac{i_A}{1 - p_B} - h_A, \frac{i_B}{1 - p_A}\right). \tag{20}$$

If this inequality is not satisfied, the fine is so large that expected losses for Alice or Bob from sharing exceed their expected gains from the exchange of the content.

6.4. Condition 3

Alice should prefer using SPIES it to simply selling the content without incentives to stop unauthorized sharing. Since Alice is the content provider, it is her choice whether to use a rights-management scheme or not.

We assume that Conditions 1 and 2 hold. Alice's expected utility from using SPIES is given by

$$i_A - (1 - p_B)(f + h_A), \tag{21}$$

her expected utility from cooperating.

Alice can assume that if she does not use SPIES, Bob will certainly share the content; so her utility if she simply sells the content without using SPIES is

$$i_A - h_A. \tag{22}$$

Alice will only use SPIES if it is superior to simply selling the content, i.e. if

$$i_A - (1 - p_B)(f + h_A) > i_A - h_A. \tag{23}$$

Solved for the fine f, this inequality reads

$$f < \frac{h_A p_B}{1 - p_B}. \tag{24}$$

If Alice is not harmed much by selling the content, then she will have to set the fine very low to gain any benefit from SPIES.

When Conditions 1, 2, and 3 are met, Inequalities 15, 20, and 24 are satisfied. Combining these, we find that SPIES can be effectively used when

$$\max\left(\frac{s}{p_B}, \frac{s}{p_A}\right) < f < \min\left(\frac{i_A}{1 - p_B} - h_A, \frac{i_B}{1 - p_A}, \frac{h_A p_B}{1 - p_B}\right) \tag{25}$$

and

$$b < \min\left(\frac{2p_B f - 2s}{1 + p_B}, \frac{2p_A f - 2s}{1 + p_A}\right) \tag{26}$$

are satisfied, which is possible when

$$\max\left(\frac{s}{p_B}, \frac{s}{p_A}\right) < \min\left(\frac{i_A}{1 - p_B} - h_A, \frac{i_B}{1 - p_A}, \frac{h_A p_B}{1 - p_B}\right). \tag{27}$$

7. Variations

In this section, we discuss modifications to SPIES to handle users who aren't trusted to pay a fine, to allow more than two legitimate possessors of the content, and to protect content known to just one person.

7.1. Security Deposits

SPIES requires that users commit to paying a fine, either through legal means (for example, a digitally signed legal document) or through technological means (for example, giving the Escrow service digital cash). In some situations it will not be possible to enforce such a commitment. For instance, in the case of a legal document, an individual may not be able to be found to pay the fee. Alternately, the machinery to support such commitments may simply not be available. In these cases, a security deposit can be used instead of a commitment to pay a fine.

Instead of sending a commitment to pay a fine during the content exchange, Alice and Bob send a security deposit. Rather than obtaining a fine from Alice and Bob if the content is shared, the Escrow service simply pays the bounty out of Alice and Bob's security deposits. At the end of the Protection Period, Alice and Bob's security deposits are returned.

Unfortunately, the use of security deposits requires that the Protection Period be finite, which limits the applicability of SPIES to content which only has value for a short period.

7.2. More than two users

SPIES can be used with multiple authorized possessors with minor modifications, but if there are more than two legitimate possessors, then *all* players must cooperate to prevent the fine from being levied. If the set of authorized possessors is X, then the expected utility for a player $x \in X$ is:

$$u_x(C) = i_x - (1 - \prod_{u \in X, u \neq x} p_u)f, \tag{28}$$

where p_u is x's estimate of u's probability of cooperating.

If x estimates a single probability p for the likelihood of cooperation for each other player in X, she estimates a return from cooperating of

$$u_X(C) = i_X - (1 - p^{|X|-1})f. \tag{29}$$

She will only participate if she expects a positive return, so for an estimated probability p of cooperation, at most

$$|X| = \log_p(1 - \frac{i_X}{f}) + 1 \tag{30}$$

legitimate possessors can participate before x chooses not to.

Assuming that the content provider has set the fine and bounty payments correctly, x should be unlikely to share the content. She will simply not participate if it is not in her interest.

As discussed in Section 4.3, watermarking can be used to reduce the multi-player case to a two player case.

7.3. Single-Player SPIES

One way of fostering cooperating in Assurance games is to change the single stage assurance game into a two-stage game where Alice observably moves first and Bob moves second. Since Bob does not have to worry about Alice defecting, no vicious circle results. Bob can evaluate his options without worrying about Alice; if he cooperates, he is guaranteed not to have to pay the fine. If he shares the content, he will have to pay the fine. Since she is known not to have shared the content, Alice need not pay the fine if the content is shared, so it is in Bob's interest to cooperate so long as $f - s > b$, the loss from sharing exceeds the bounty payment.

It is possible to achieve this in SPIES when only Bob, not Alice, knows the protected content. In this case, Alice is actually unable to defect and register or share the content. This occurs, for example, in a secure key exchange protocol like JFK (Aiello et al., 2002). Bob can use a cut and choose protocol to commit to his secret with the Escrow service, and can then be certain that Alice will not share his secret. For her part, Alice can be certain that she will be notified, and that Bob will have to pay a fine, if he shares the secret. Asymmetric keys used as passwords can also be protected in this way.

8. Conclusion

For rights-protected content which is not widely shared and satisfies certain economic requirements, SPIES can be used to both discourage and detect sharing. The protocol does not rely on any traditional DRM methods but can be used in conjunction with them. In the scheme, users agree to pay a fee if protected content is shared. Anyone who has the content can get a bounty by registering.

A game-theoretic analysis of SPIES shows that rational participants will only share the content in unauthorized ways when they fear that the other participant will share. Several techniques can be used to reduce this mistrust between participants. In addition, sharing is quickly discovered, because it is in the interest of illegitimate sharers to claim the bounty before sharing.

SPIES works properly only when the fine and bounty payment are set appropriately by the content provider. For some content-protection problems, it is not be possible to set the fine and bounty so that they satisfy all of the requirements simultaneously. The economic conditions required for SPIES can be reduced to a necessary and sufficient inequality involving the the content provider and consumer's valuations of the content, their trust in one other, the content-provider's level of aversion to illegitimate sharing, and the total market value of the content.

Many variations of SPIES are possible. One variation enables it to use security deposits rather than fines. Sharing content with more than two legitimate possessors is possible when trust is high; watermarking can be used to reduce the effective number of possessors to two. It is also possible to share content without the content provider being aware of what is shared; in this case, the content consumer can have perfect trust that the content provider will not share the content pre-emptively.

SPIES is applicable to passwords, corporate knowledge-sharing, pre-release media, and many other types of content. Its chief limitations are the requirements for a limited number of legitimate possessors and that it provides incentives rather than provably preventing sharing. Since technological and legal attempts to prevent unauthorized sharing have mostly had the effect of erecting economic barriers to sharing, we believe that the use of direct economic incentives to limit sharing is a useful addition to traditional DRM techniques.

References

Aiello, W., Bellovin, S. M., Blaze, M., Ioannidis, J., Reingold, O., Canetti, R., Keromytis, A. D., 2002. Efficient, DoS-resistant, secure key exchange for Internet protocols. In: CCS 2002: Proceedings of the 9th ACM conference on Computer and Communications Security. ACM Press, New York, NY, USA, pp. 48–58.

Arvan, L., Cabral, L., Santos, V., 1999. Meaningful cheap talk must improve equilibrium payoffs. Mathematical Social Sciences 37, 97–106.

Baliga, S., Sjöström, T., 2004. Arms races and negotiations. Review of Economic Studies 71, 351–369.

Bao, F., Deng, R. H., Mao, W., May 1998. Efficient and practical fair exchange protocols with off-line TTP. In: Proc. 1998 Ieee Symposium On Security And Privacy. IEEE Computer Society Press, Oakland, California, USA.

Berthold, O., Federrath, H., Kohntopp, M., April 2000. Project anonymity and unobservability in the Internet. In: Computers Freedom And Privacy Conference (CFP 2000) Workshop on Freedom and Privacy by Design. ACM Press, Toronto, Canada.

Boneh, D., Shaw, J., 1995. Collusion-secure fingerprinting for digital data. Lecture Notes in Computer Science 963, 452–465.

Chaum, D., Crépeau, C., Damgard, I., 1988. Multiparty unconditionally secure protocols. In: STOC '88: Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing. ACM Press, New York, NY, USA, pp. 11–19.

Cox, I. J., Kilian, J., Leighton, T., Shamoon, T., 1996. A secure, robust watermark for multimedia. In: Information Hiding: First International Workshop. Springer-Verlag, Cambridge, UK, pp. 187–206.

Dittmann, J., Behr, A., Stabenau, M., Schmitt, P., Schwenk, J., Ueberberg, J., 2000. Combining digital watermarks and collusion secure fingerprints for digital images. SPIE Journal of Electronic Imaging 9, 456–467.

Douceur, J. R., March 2002. The Sybil attack. In: Proc. of the IPTPS02 Workshop. Springer-Verlag, Cambridge, MA, USA.

Farrell, J., 1984. Cheap talk, coordination, and entry. The RAND Journal of Economics 18 (1), 34–39.

Freedman, M. J., Morris, R., November 2002. Tarzan: A peer-to-peer anonymizing network layer. In: Proc. ACM Conference on Computer and Communications Security (CCS 2002). ACM Press, Washington, DC, USA.

Gibbons, R., 1992. Game Theory for Applied Economists. Princeton University Press, Princeton, NJ, USA.

Golle, P., Leyton-Brown, K., Mironov, I., 2001. Incentives for sharing in peer-to-peer networks. In: EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce. ACM Press, New York, NY, USA, pp. 264–267.

Haber, S., Horne, B., Pato, J., Sander, T., Tarjan, R. E., 2003. If piracy is the problem, is DRM the answer? In: Becker, E., Buhse, W., Günnewig, D., Rump, N. (Eds.), Digital Rights Management. Vol. 2770. Springer-Verlag, Berlin, Germany, pp. 224–233.

Horne, B., Pinkas, B., Sander, T., 2001. Escrow services and incentives in peer-to-peer networks. In: Proceedings of the 3rd ACM Conference on Electronic Commerce. ACM Press, Tampa, Florida.

Linnartz, J.-P. M. G., 1998. The "ticket" concept for copy control based on embedded signalling. Lecture Notes in Computer Science 1485, 257–274.

Margolin, N. B., Wright, M., Levine, B. N., Oct. 2004a. Analysis of an incentives-based protection system. In: Proc. ACM Digital Rights Management Workshop. ACM Press, Washington, DC, USA.

Margolin, N. B., Wright, M., Levine, B. N., June 2004b. SPIES: Secret Protection Incentive-based Escrow System. In: Workshop on the Economics of Peer-to-Peer Systems (P2PEcon).

Reiter, M. K., Rubin, A. D., November 1998. Crowds: Anonymity for Web Transactions. ACM Transactions on Information and System Security 1 (1), 66–92.

Schneier, B., 1996. Applied Cryptography. John Wiley & Sons, New York, NY, USA.

Sibert, O., Bernstein, D., Wie, D. V., 1995. The DigiBox: A self-protecting container for information commerce. In: Proceedings of the First USENIX Workshop on Electronic Commerce. USENIX Association, New York, NY, USA, pp. 171–183.

Syverson, P. F., Goldschlag, D. M., Reed, M. G., 4–7 1997. Anonymous connections and onion routing. In: IEEE Symposium on Security and Privacy. IEEE Computer Society Press, Oakland, California, pp. 44–54.

A. Game Theory

A.1. Basic Concepts

A *game* in game theory consists of several players, each with different strategy choices and preferences, and the outcomes for all given the different possible combinations of choices.

For instance, consider the *Rock-Paper-Scissors* game. Two players, 1 and 2, simultaneously choose either Rock, Paper, or Scissors. Among these choices, Scissors dominates Paper, Paper dominates Rock, and Rock dominates Scissors. The player with the dominating choice wins. If the choices are the same, the game is a draw. The utility of winning is 1, the utility of losing is 0, and the utility of a draw is 0.5.

In this game, the players are 1 and 2:

$$P = \{1, 2\}.$$

The strategies are the same for both players, and consist of the set

$$\sigma = \{\mathit{Rock}, \mathit{Paper}, \mathit{Scissors}\}.$$

The utility functions for the players 1 and 2 are given by

$$\begin{aligned} u_1(\sigma_1, \sigma_2) &= \begin{cases} 1 & \text{if } \sigma_1 \text{ dominates } \sigma_2 \\ 0.5 & \text{if } \sigma_1 = \sigma_2 \\ 0 & \text{if } \sigma_2 \text{ dominates } \sigma_1 \end{cases} \\ u_2(\sigma_1, \sigma_2) &= u_1(\sigma_2, \sigma_1) \end{aligned}$$

These two utility functions are represented by a single matrix, where by convention Player 1's strategies are given by the rows and Player 2's strategies by the columns. The matrix entries give the the utilities of the players as an ordered pair of Player 1's utility and Player 2's utility.

	Rock	*Paper*	*Scissors*
Rock	(0.5,0.5)	(0,1)	(1,0)
Paper	(1, 0)	(0.5,0.5)	(0,1)
Scissors	(0, 1)	(1,0)	(0.5,0.5)

In this paper, player's utilities are real numbers representing units of currency. Utility functions map sets of strategies to sets of utilities:

$$u : \sigma^n \to \mathbb{R}^n,$$

where n is the number of players. By definition, rational players try to maximize the dimension of u representing their own utility.

A central concept in game theory is the *Nash Equilibrium*, a state in which no player can profitably change her strategy, given the other player's choices of strategy. An outcome (a choice of strategies for each player, or a choice of probabilities over the strategy space in mixed-strategy games) is *Pareto optimal* if there is no state that is better for some player and at least as good for all others. In the *Rock-Paper-Scissors* game, there are no Nash Equilibria: for each outcome, the losing player would rather change her choice to one that dominates the other player's choice.

A strategy X *strictly dominates* a strategy Y for a player if, no matter what strategies the other players choose, playing X leads to a greater utility for that player than playing Y. A strategy X *weakly dominates* a strategy Y for a player if playing X leads to at least as great utility for that player than playing Y no matter what the other players' chosen strategies are, and there is *some* choice of strategies the other players can make for which playing X leads to a strictly greater utility than playing Y.

A.2. Prisoner's Dilemma and Assurance Games

In this section, we discuss two well-known games with applications to the SPIES protocol.

The prisoner's dilemma is one of the best known games in game theory. Suppose that Player 1 and Player 2 are criminal accomplices captured by the police. If neither confesses to their crime, the police will convict them both of minor offenses, and both go to jail for a year. If just one confesses, that prisoner is set free as a reward; the police then have enough evidence to send the other prisoner to jail for 20 years. If both prisoners confess, both receive a five year sentence. Staying silent in this game is called *Cooperating*, and confessing is called *Defecting*.

The Prisoner's Dilemma is represented by the following matrix.

	Cooperate	*Defect*
Cooperate	(-1,-1)	(-20,0)
Defect	(0, -20)	(-5,-5)

The game has a single Nash equilibrium at $(\mathit{Defect}, \mathit{Defect})$: No matter what her opponent's choice, a prisoner is better off confessing than remaining silent. Surprisingly, the rational outcome produced by the individual reasoning of the two players, (*Defect*, *Defect*), is inferior for both players to (*Cooperate*,*Cooperate*).

The *Assurance Game* is a game closely related to the Prisoner's Dilemma game. While players in the Prisoner's dilemma game always prefer defecting to cooperating – even though it leads to worse outcomes for both – the players in the Assurance game generally prefer to cooperate. They only choose to defect if they fear that the other player will also defect. The traditional example of an Assurance game is a stag hunt. One hunter scares the stag towards the hiding place for the other; unless both participate, the stag escapes. Both hunters want to catch the stag, but they don't want to waste their time. If one hunter fears the other will not show up, he may also stay home.

The strategies and utilities of the Assurance game are as follows: This set of utilities is represented as follows.

	Cooperate	*Defect*
Cooperate	(10, 10)	(1,5)
Defect	(5, 1)	(3,3)

The key difference is that in this case, playing *Defect* does not always lead to higher utility. If the opponent has played *Defect*, then *Defect* is the best choice, but if she has played *Cooperate*, then *Cooperate* is the best choice. Therefore the Assurance game has two Nash equilibria, one at (*Cooperate*,*Cooperate*) and one at (*Defect*,*Defect*). The first of these equilibria has a higher utility for both players. However, the strategy *Defect* is less risky: a player will get a utility of at least 3 by playing *Defect*, as opposed to a possibility of getting a utility 1 if *Cooperate* is played.

If both players know that each other's outcome preferences are perfectly rational, and both have full knowledge of this (and know this fact, and the fact that they both know it, and so on, a condition known as *full information*) then they will both play *Cooperate*. However, if there is any doubt, as in many real-world situations, there is the potential for a vicious circle. Player 1 may reason that there is at least some chance that Player 2 is irrational and will defect — so perhaps she should defect pre-emptively. Player 2 can see that Player 1 may reason that way and defect — so perhaps she should too, even if she is certain that Player 1 is rational. Such reasoning only increases Player 1's uncertainty about Player 2's choices, and there is no end to the cycle. For this reason there is a chance in real-world situations that even rational players will defect, and this chance is not easily quantifiable.

There are three methods for increasing the likelihood of the cooperative outcome in the assurance game.

1. Change the game so that the parties can openly commit to *Cooperate*, eliminating doubts that can lead to a vicious cycle of uncertainty and defection.
2. Reduce the penalty for playing *Cooperate* when the opponent *Defects*, for example by lowering the defector's utility in (*Defect*, *Cooperate*) and (*Cooperate*,*Defect*) to 3 (so that the utilities of these states are (3,1) and (1,3) respectively), or by raising the co-operator's utility so that these state are (5,3) and (3,5), respectively. Playing *Defect* becomes less likely since it offers only a marginal improvement of utility over *Cooperate* if the opponent plays *Defect*.
3. Improve trust between the participants to reduce the possibility of a vicious circle effect.

The second and third of these techniques were behind the Mutually Assured Destruction (MAD) strategy that the United States and the Soviet Union employed during the Cold War (Baliga and Sjöström, 2004). The arms race the two countries were engaged in can be modeled as an Assurance game, a nuclear attack being *Defection*.

The use of a large number of powerful nuclear warheads ensured that any improved utility a country obtained from making preemptive attack would be purely academic. So distrust would have to be almost total before a preemptive attack would appear attractive.

The countries also stayed in contact with each other at a high level. This, and the well-known high cost of a nuclear war, improved each country's confidence that the rulers of the other preferred peace to nuclear war.