

# DISCUSSION PAPER SERIES

DP17865

## **DIGITAL HERMITS**

Jeanine Miklos-Thal, Avi Goldfarb, Avery Haviv and  
Catherine Tucker

**INDUSTRIAL ORGANIZATION**

**CEPR**

# DIGITAL HERMITS

*Jeanine Miklos-Thal, Avi Goldfarb, Avery Haviv and Catherine Tucker*

Discussion Paper DP17865  
Published 31 January 2023  
Submitted 31 January 2023

Centre for Economic Policy Research  
33 Great Sutton Street, London EC1V 0DX, UK  
Tel: +44 (0)20 7183 8801  
[www.cepr.org](http://www.cepr.org)

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Jeanine Miklos-Thal, Avi Goldfarb, Avery Haviv and Catherine Tucker

# DIGITAL HERMITS

## Abstract

When a user shares multi-dimensional data about themselves with a firm, the firm learns about the correlations of different dimensions of user data. We incorporate this type of learning into a model of a data market in which a firm acquires data from users with privacy concerns. User data is multi-dimensional, and each user can share no data, only non-sensitive data, or their full data with the firm. As the firm collects more data and becomes better at drawing inferences about a user's privacy-sensitive data from their non-sensitive data, the share of new users who share no data ("digital hermits") grows. At the same time, the share of new users who share their full data also grows. The model therefore predicts a polarization of users' data sharing choices away from non-sensitive data sharing to no sharing and full sharing.

JEL Classification: D42, D82, D83, L20

Keywords: Digital markets

Jeanine Miklos-Thal - [jeanine.miklos-thal@simon.rochester.edu](mailto:jeanine.miklos-thal@simon.rochester.edu)  
*University of Rochester Simon School and CEPR*

Avi Goldfarb - [Avi.Goldfarb@Rotman.Utoronto.Ca](mailto:Avi.Goldfarb@Rotman.Utoronto.Ca)  
*University of Toronto, Rotman School of Management*

Avery Haviv - [avery.haviv@simon.rochester.edu](mailto:avery.haviv@simon.rochester.edu)  
*University of Rochester Simon School*

Catherine Tucker - [cetucker@mit.edu](mailto:cetucker@mit.edu)  
*MIT Sloan School Of Management*

# Digital Hermits

Jeanine Miklós-Thal      Avi Goldfarb      Avery Haviv      Catherine Tucker

January 27, 2023

## Abstract

When a user shares multi-dimensional data about themselves with a firm, the firm learns about the correlations of different dimensions of user data. We incorporate this type of learning into a model of a data market in which a firm acquires data from users with privacy concerns. User data is multi-dimensional, and each user can share no data, only non-sensitive data, or their full data with the firm. As the firm collects more data and becomes better at drawing inferences about a user’s privacy-sensitive data from their non-sensitive data, the share of new users who share no data (“digital hermits”) grows. At the same time, the share of new users who share their full data also grows. The model therefore predicts a polarization of users’ data sharing choices away from non-sensitive data sharing to no sharing and full sharing.

---

Jeanine Miklós-Thal is Professor of Economics and Management at Simon Business School, University of Rochester and Research Fellow at CEPR; e-mail: [jeanine.miklos-thal@simon.rochester.edu](mailto:jeanine.miklos-thal@simon.rochester.edu). Avi Goldfarb is Professor of AI and Healthcare at the Rotman School of Management, University of Toronto and Research Associate at NBER. Avery Haviv is Associate Professor of Marketing at the Simon Business School, University of Rochester. Catherine Tucker is the Sloan Distinguished Professor of Management at MIT and Research Associate at NBER. We thank seminar and conference participants at the 2022 MaCCI Summer Institute in Competition Policy and USTC-UIUC. We also thank Zheng Gong for excellent research assistance. All mistakes are our own.

# 1 Introduction

Users consider some data more privacy sensitive than other data (Acquisti et al., 2016). For example, a consumer might be unconcerned if a company has access to data on their favorite flower, but very concerned if a company has access to data that reflected their sexual orientation. Traditionally, privacy preferences and data-sharing decisions across users have reflected these heterogeneous preferences over which types of data are sensitive and which types of data are innocuous (Prince and Wallsten, 2022).

However, this division of data into deemed ‘sensitive’ and deemed ‘not sensitive’ is challenged by advances in machine learning and prediction. This is because advances in prediction and inference mean that potentially sensitive data can be inferred from data that otherwise seems innocuous. For example, it might be possible to learn the likelihood of someone having a particular sexual orientation from their flower-buying preferences and other seemingly innocuous data about them. Though this may seem far-fetched, existing research has shown that traits such as IQ can be predicted from someone liking curly fries on Facebook, or sexual orientation predicted from liking the Wu-Tang Clan on Facebook (Kosinski et al., 2013). This leads to a new way of viewing privacy preferences, where rather than individuals being concerned about the sensitivity of a piece of data, they are instead concerned about what information about them can be inferred from a piece of data, because that data is correlated with other potentially more privacy sensitive traits.

This paper analyzes the implications for users’ data-sharing decisions of such privacy preferences. We build a theoretical model in which data about each user is multi-dimensional and varying across dimensions in the extent to which it is privacy sensitive. Users can decide to share all their data, no data, or only a subset of their data. New users enter this digital economy gradually over time. A firm uses accumulated data and machine learning to learn about the correlations between the different dimensions of user data. The potential for correlation between data that seems innocuous and data that is privacy sensitive motivates our setup.

Our analysis shows that as a firm accumulates more data and predicts more accurately the correlation between different dimensions of user data, users’ data sharing decisions polarize. First, more users choose to share no data. This group are “digital hermits.” Second, more users choose to share all their data, whether sensitive or not sensitive. The share of digital hermits and the share

of full data sharers in the population therefore both rise over time, while the share of users who share only non-sensitive data falls over time.

This polarization occurs because extensive data sharing by early users imposes a negative externality on later users who share only non-sensitive data. This externality results from the firm’s learning about the correlation between different dimensions of user data. More data sharing by early users enables the firm to draw more accurate inferences about a later user’s sensitive data from that user’s non-sensitive data. This pushes up the compensation required to encourage non-sensitive data sharing by users who are concerned about the use of this data to infer, through correlation, sensitive information. Later users who value privacy sufficiently choose to be digital hermits who share no data at all, even data that they do not consider privacy sensitive and would be happy to share with the firm in the absence of the learning-induced externality. Later users who value privacy less, realizing that there is little point in keeping just some data private, share all their data, including sensitive data that they would have preferred to keep private in the absence of the firm’s learning.

The polarization in data sharing decisions and the associated rise of digital hermits occur both when the benefits that users receive for sharing data are fixed over time and when the firm dynamically adjusts the benefits of sharing data—or data prices—offered to users. In the latter case, as the firm learns about the correlation between different data dimensions, the share of users who choose to be digital hermits in the population rises, although the firm raises both the price for non-sensitive data and the price for a user’s full data.

This paper also shows that a more patient firm optimally commits to collecting data more slowly. A forward-looking firm anticipates that as its ability to predict privacy-sensitive data from non-sensitive data improves, future users will demand higher compensation to share non-sensitive data. Inducing current users to share all their data therefore hurts the firm’s future profits. Because of this, a more patient firm offers a lower price for the privacy-sensitive data in early periods in order to slow down its own learning and delay the erosion in future profits due to users’ privacy concerns.

This paper contributes to the recent theory literature on the consequences of learning and data externalities across individuals in digital markets (Choi et al., 2019; Ichihashi, 2021; Bergemann et al., 2022; Acemoglu et al., 2022). The main modeling differences lie in the multi-dimensionality of

each user’s data and the focus on the consequences of a firm learning the correlation patterns across these different data dimensions. This learning about correlation provides a new microfoundation for the externalities across users that arise in data markets, which were assumed in Choi et al. (2019), Ichihashi (2021), and Acemoglu et al. (2022).

These modeling differences lead our model to have some starkly different predictions. The existing literature emphasizes that learning and data externalities across users often lead to lower prices for data and excessive data sharing in equilibrium. While our finding that more users share all data as the firm becomes better at making predictions mirrors the excessive data sharing result in recent literature, we also obtain the sharply contrasting insight that more users will shut down data sharing completely and choose to be digital hermits. Moreover, the learning-induced data externalities in our model raise, rather than reduce, the equilibrium price for non-sensitive data as well as the total price for a user’s full data.

## 2 A model of data markets with multi-dimensional user data

**Setting** Time is discrete and indexed by  $t \in \{0, 1, 2, \dots\}$ . There is one firm and infinitely many users  $C_1, C_2, \dots$ , with user  $C_t$  being alive in period  $t$ . We assume users are alive for just one period to abstract away from dynamic considerations of user data today affecting what the firm can learn about that same user in the future, a point which is well-addressed by Ichihashi (2020). User  $C_t$ ’s type is a vector  $(x_t, y_t)$  consisting of the realizations of two random variables  $X_t$  and  $Y_t$ . For every  $t \geq 1$ , the bivariate random vector  $(X_t, Y_t)$  is distributed according to the bivariate probability mass function

$$f(x, y|\rho) = \begin{cases} \frac{1+\rho}{4} & \text{if } (x, y) \in \{(0, 0), (2\sigma_X, 2\sigma_Y)\} \\ \frac{1-\rho}{4} & \text{if } (x, y) \in \{(0, 2\sigma_Y), (2\sigma_X, 0)\} \end{cases},$$

where  $\rho$  is the correlation of  $X_t$  and  $Y_t$ . If  $\rho > 0$ , the high (low) realization of  $X_t$  is associated with the high (low) realization of  $Y_t$ ; if  $\rho < 0$ , the low (high) realization of  $X_t$  is associated with the high (low) realization of  $Y_t$ . Conditional on  $\rho$ ,  $(X_t, Y_t)$  and  $(X_s, Y_s)$  are mutually independent for any  $t$  and  $s \neq t$ .

The firm learns about the correlation  $\rho$  from the data shared by users. Formally, we assume that at the beginning of the game, nature draws  $\rho \in \{-r, r\}$ , with each of the two realizations

being equally likely and  $r \in (0, 1)$ . Hence, the firm’s prior expectation is  $E[\rho] = \frac{r}{2} + \frac{-r}{2} = 0$ . As the firm accumulates data, it uses Bayesian updating to revise its belief about  $\rho$ .

The marginal probability mass functions  $f_X(x) = \sum_y f(x, y|\rho) = \frac{1}{2}$  and  $f_Y(y) = \sum_x f(x, y|\rho) = \frac{1}{2}$  are independent of  $\rho$  in our setting. Hence, as the firm learns about  $\rho$ , it learns about the joint probability function  $f(x, y|\rho)$ , but not about the marginal probability functions. This reflects that in a world with multi-dimensional user characteristics, even if the proportion of users with a particular characteristic (say,  $x_t = 0$ ) in the population is known from aggregate data, data analytics can still be useful to learn about the correlation patterns between different characteristics.

User  $C_t$  possesses a data vector  $(d_{X_t}, d_{Y_t})$  that, if shared with the firm, allows the firms to draw inferences about  $(x_t, y_t)$ . To focus on the firm’s learning about the correlation  $\rho$ , we assume that the firm’s inference from data to type is perfect within each of the two dimensions. The firm can perfectly infer  $x_t$  from  $d_{X_t}$  and  $y_t$  from  $d_{Y_t}$ . Importantly, a user is able to share only a subset of their data, for example only  $d_{X_t}$ , with the firm. In this case, the firm uses its posterior belief about  $\rho$  to draw inferences about  $C_t$ ’s type in the dimension about which  $C_t$  did *not* share data from its knowledge of  $C_t$ ’s type in the dimension about which  $C_t$  did share data.

**Timing** The game proceeds as follows. In period  $t = 0$ , nature draws the correlation  $\rho \in \{-r, r\}$ . In every period  $t \geq 1$ , the firm offers user  $C_t$  a price  $p_{X_t} \geq 0$  for sharing data  $d_{X_t}$  and a price  $p_{Y_t} \geq 0$  for sharing data  $d_{Y_t}$ .<sup>1</sup> Though we use the term ‘price’, we recognize that formal markets for data rarely exist, so this price may reflect other benefits that a user receives from exchanging that particular piece of data with the firm.

Next, each user  $C_t$  decides whether to share no data, only data  $d_{X_t}$ , or their full data vector  $(d_{X_t}, d_{Y_t})$  with the firm. We will assume below that data dimension  $Y$  is privacy-sensitive but dimension  $X$  is not. This implies that sharing the full data vector dominates sharing only  $d_{Y_t}$  for every  $C_t$ . Therefore, we ignore the potential for sharing the sensitive data of dimension  $Y$  only, without losing generality.  $C_t$ ’s data sharing decision is denoted by  $s_t \in \{N, I, F\}$ , where  $N$  stands for sharing no data,  $I$  for sharing non-sensitive (or ‘innocuous’) data, and  $F$  for sharing full data. Finally, the firm analyzes the data shared by  $C_1$  to  $C_t$  to predict  $(x_t, y_t)$  and period  $t$  payoffs are earned.

---

<sup>1</sup>Allowing the firm to offer a third price for a user’s full data vector (i.e., the bundle of  $d_{X_t}$  and  $d_{Y_t}$ ) would not affect the results.



**Payoffs** For any  $t \geq 1$ , let  $\mathbf{D}_t = (D_1, D_2, \dots, D_t)$  denote the vector of data shared by  $C_1$  to  $C_t$ , where  $D_t = (\phi, \phi)$  if  $s_t = N$ ,  $D_t = (d_{Xt}, \phi)$  if  $s_t = I$ , and  $D_t = (d_{Xt}, d_{Yt})$  if  $s_t = F$ .

The firm derives revenues from predicting users' types, for instance, because it can better persuade users to make purchases if it knows more about them. Specifically, following Acemoglu et al. (2022), we assume that the firm's payoff is decreasing in the mean squared errors of the firm's estimators of user types. The key difference to Acemoglu et al. (2022) is that, in our setting, each user's type is multi-dimensional. Gross of any payments made to  $C_t$ , the firm's current-period payoff in period  $t \geq 1$  is

$$\frac{\sigma_X^2 - E \left[ (\hat{x}_t(\mathbf{D}_t) - X_t)^2 \right]}{\sigma_X^2} + \frac{\sigma_Y^2 - E \left[ (\hat{y}_t(\mathbf{D}_t) - Y_t)^2 \right]}{\sigma_Y^2},$$

where  $\hat{x}_t(\mathbf{D}_t)$  is the firm's Bayes estimator of  $x_t$  as a function of the data observed by the firm, and  $\hat{y}_t(\mathbf{D}_t)$  is the firm's Bayes estimator of  $y_t$  as a function of the data observed by the firm. The variance terms are included as convenient normalizations, ensuring that each of the two terms in the firm's payoff function lies between 0 and 1.

Users value privacy in dimension  $Y$ , which could be for intrinsic reasons or for instrumental reasons. Again following Acemoglu et al. (2022), we model consumer privacy concerns in a reduced-form manner using the same mean-squared error of the firm's estimator in the privacy-sensitive dimension  $Y$  as in the firm's profit function. Gross of any payments from the firm,  $C_t$ 's payoff is given by

$$v_t \frac{E \left[ (\hat{y}_t(\mathbf{D}_t) - Y_t)^2 \right] - \sigma_Y^2}{\sigma_Y^2},$$

where  $v_t \geq 0$  measures  $C_t$ 's privacy valuation. Each user's payoff is thus increasing in the mean squared error of the firm's estimator in the privacy-sensitive dimension in the period in which the user is alive.

The users' privacy valuations  $(v_0, v_1, v_2, \dots)$  are i.i.d. draws from a uniform distribution on  $[0, V]$ , where  $V > 2$ . The CDF of this distribution is denoted by  $H(v) = \min \left\{ \frac{v}{V}, 1 \right\}$  for  $v \geq 0$ . The privacy valuation  $v_t$  is  $C_t$ 's private information.

An important assumption in our model is that  $C_t$ 's payoff is independent of the user's realized type  $(x_t, y_t)$ . One interpretation of this assumption is that users have an intrinsic preference for

keeping information about dimension  $Y$  private, regardless of their actual type. This could be the case, for example, if the information concerns a user’s sexual activity, political views, psychological traits, or another category deemed intrinsically sensitive by users.

Another possible interpretation is that, although a user would want to keep only information about a specific realization of  $Y_t$  private, the user is not able to infer  $(x_t, y_t)$  from their own personal data  $(d_{X_t}, d_{Y_t})$ , unlike the firm. This could be the case, for instance, if a user’s type concerns personal health risks that are unknown to the user, but that can potentially be predicted by data-rich firms.

A final interpretation is that users fear that certain type realizations may have grave repercussions when revealed or inferred, but, at the time of the user’s data sharing decision, there is uncertainty about which realization will have negative repercussions, as in the case of allegiance to a particular political leader when it is uncertain who will win an election. This again creates an incentive for users to keep their data private regardless of their realized type.

The assumption that each user’s payoff is independent of their realized type prevents situations in which the firm draws inferences about a user’s type from the user’s decision to withhold data. This rules out the classic “unraveling” result in information disclosure models, according to which all users reveal their types without receiving any compensation or where markets fail to enable Pareto-improving exchange (Grossman, 1981; Grossman and Hart, 1980; Milgrom, 1981; Taylor, 2004; Acquisti and Varian, 2005). Related recent theoretical work on externalities in data markets makes similar assumptions. Bergemann et al. (2022) assume that consumers and the data intermediary in their model hold symmetric information at the data contracting stage. Acemoglu et al. (2022), in whose model a user’s data consists of a signal about their one-dimensional type, assume that (conditional on the firm’s forecast) a user’s payoff is independent of the user’s signal about their own type. Choi et al. (2019) do not analyze an explicit statistical model with belief updating. Instead, they assume that when buying a service from the firm, which requires data sharing, the consumer pays a nuisance cost that is increasing in the total number of consumers sharing their data.

**Prediction quality** Given the assumption that there is no learning about the marginal probability mass functions of  $X_t$  and  $Y_t$ , we have that  $E \left[ (\hat{x}_t(\mathbf{D}_t) - X_t)^2 \right] = \sigma_X^2$  and  $E \left[ (\hat{y}_t(\mathbf{D}_t) - Y_t)^2 \right] =$

$\sigma_Y^2$  if  $s_t = N$ . Moreover, since the firm can fully infer  $C_t$ 's type in the dimension(s) in which  $C_t$  shares data,  $E \left[ (\hat{x}_t(\mathbf{D}_t) - X_t)^2 \right] = 0$  if  $s_t = I$  or  $s_t = F$  and  $E \left[ (\hat{y}_t(\mathbf{D}_t) - Y_t)^2 \right] = 0$  if  $s_t = F$ .

Finally, consider the mean-squared error of the firm's estimator  $\hat{y}_t$  when  $C_t$  shares only non-sensitive data ( $s_t = I$ ). Noting that only full data vectors (i.e., only data vectors of the form  $(d_{X_t}, d_{Y_t})$ ) contain information useful for the firm to revise its belief about  $\rho$ , we denote by  $n_{t-1} = |\{j \in \{1, \dots, t-1\} : s_j = F\}|$  the number of users who shared their full data vectors in periods 1 to  $t-1$  and by  $k_{t-1} = |\{j \in \{1, \dots, t-1\} : s_j = F, (x_j, y_j) \in \{(0, 0), (2\sigma_X, 2\sigma_Y)\}\}|$  how many of these vectors indicated positive correlation. The firm's posterior belief about the correlation  $\rho$  can then be denoted by  $\hat{\rho}(k_{t-1}|n_{t-1}) = E[\rho|k_{t-1}, n_{t-1}]$ , and the firm's estimator  $\hat{y}_t$  becomes

$$\hat{y}_t(\mathbf{D}_t) = \hat{y}_t(\hat{\rho}(k_{t-1}|n_{t-1}), x_t) = \begin{cases} (1 + \hat{\rho}(k_{t-1}|n_{t-1}))\sigma_Y & \text{if } x_t = 2\sigma_X, \\ (1 - \hat{\rho}(k_{t-1}|n_{t-1}))\sigma_Y & \text{if } x_t = 0. \end{cases}$$

The mean squared error of the firm's estimator  $\hat{y}_t$  given  $n_{t-1}$  is thus equal to

$$\begin{aligned} & E \left[ (\hat{y}_t(\hat{\rho}(k_{t-1}|n_{t-1}), x_t) - Y_t)^2 \right] \\ &= E \left[ \frac{1 + \hat{\rho}(k_{t-1}|n_{t-1})}{2} (1 - \hat{\rho}(k_{t-1}|n_{t-1}))^2 \sigma_Y^2 + \frac{1 - \hat{\rho}(k_{t-1}|n_{t-1})}{2} (1 + \hat{\rho}(k_{t-1}|n_{t-1}))^2 \sigma_Y^2 \right] \\ &= \sigma_Y^2 \left( 1 - E \left[ \hat{\rho}(k_{t-1}|n_{t-1})^2 \right] \right), \end{aligned}$$

where the expectation is taken over the possible realizations of  $k_{t-1}$ . The number of full data vectors shared by past users,  $n_{t-1}$ , can be thought of as the size of the sample available to the firm to estimate  $\rho$ , and the estimator  $\hat{\rho}(k_{t-1}|n_{t-1})$  is a random variable due to the sampling variability in  $k_{t-1}$ .

We assume that when deciding how much data to share, a user is aware of the current mean-squared error of the firm's estimator  $\hat{y}_t$  in case the user shares only non-sensitive data. Formally,  $C_t$  observes  $n_{t-1}$ . This assumption captures user awareness of the firm's current capability to draw inferences and accurately predict unshared data.

**Equilibrium concept** We analyze pure-strategy Markov Perfect equilibria, where the state is the number of past users  $n_{t-1}$  who have shared full data (which is a sufficient statistic for the mean-squared error of the firm's estimator). Strategies are assumed to be stationary in that they

depend on the state but not on the time period  $t$ . Formally, the firm's strategy is given by a pair of functions  $(P_X(n_{t-1}), P_Y(n_{t-1}))$  that specify the data prices offered given the state. And each user  $C_t$ 's strategy is given by a function  $S(n_{t-1}, p_X, p_Y; v_t)$  that specifies the user's data sharing decision  $s_t \in \{N, I, F\}$  given the state  $n_{t-1}$ , the price vector  $(p_X, p_Y)$  offered by the firm to  $C_t$ , and  $C_t$ 's privacy valuation  $v_t$ . Our generic notation for the state will be  $n$ .

In equilibrium, for every  $n \in \mathbb{N}_0$ , the firm's strategy maximizes the firm's expected discounted payoffs from the current period onwards, and the user strategy maximizes each user's payoff in the period in which they are alive. The firm's discount factor is denoted by  $\delta \in [0, 1)$ . User indifferences are broken in favor of sharing more data.

### 3 Analysis

#### 3.1 Preliminaries: Information leakage from non-sensitive data

We define the information leakage from non-sensitive data as the (normalized) reduction in the mean squared error of the firm's estimator  $\hat{y}_t$  when  $C_t$  reveals non-sensitive data versus no data:

**Definition 1** *The information leakage from non-sensitive data is*

$$L(n_{t-1}) \equiv \frac{\sigma_Y^2 - E\left[(\hat{y}_t(\hat{\rho}(k_{t-1}|n_{t-1}), x_t) - Y_t)^2\right]}{\sigma_Y^2} = E\left[\hat{\rho}(k_{t-1}|n_{t-1})^2\right].$$

Our first lemma shows that the information leakage from non-sensitive data is increasing in  $n$  at a decreasing rate.

**Lemma 1** *For every  $n \in \mathbb{N}_0$ ,*

$$L(n+1) - L(n) > L(n+2) - L(n+1) > 0,$$

*with  $L(0) = E[\rho] = 0$  and  $\lim_{n \rightarrow \infty} L(n) = E[\rho^2] = r^2$ .*

As the firm accumulates more full data vectors, the mean-squared error of its estimator of unshared sensitive data from non-sensitive data falls, and hence the information leakage from non-sensitive data rises. Moreover, the marginal effect of more data is decreasing. Figure 1 offers an illustration.




Figure 1: Information leakage from non-sensitive data ( $r = 0.5$ )

### 3.2 User data sharing decisions given prices

Consider the data sharing decision of user  $C_t$  if the current level of information leakage from non-sensitive data is  $L \in [0, 1)$  and the firm offers the data prices  $(p_X, p_Y)$ .  $C_t$ 's payoff is

$$U_F(p_X, p_Y; v_t) = p_X + p_Y - v_t$$

if they share full data,

$$U_I(p_X, p_Y, L; v_t) = p_X - v_t L$$

if they share non-sensitive data, and 0 if they share no data. Sharing full data is optimal if  $U_F \geq \max\{U_I, 0\}$ , which holds if and only if  $C_t$ 's privacy valuation

$$v_t \leq \underline{v}(L, p_X, p_Y) \equiv \min \left\{ \frac{p_Y}{1-L}, p_X + p_Y \right\}.$$

Sharing no data is optimal if  $0 > \max\{U_I, U_F\}$ , which holds if and only if  $C_t$ 's privacy valuation

$$v_t > \bar{v}(L, p_X, p_Y) \equiv \begin{cases} \infty & \text{if } L = 0 \\ \max \left\{ \frac{p_X}{L}, p_X + p_Y \right\} & \text{if } L > 0 \end{cases}.$$

If  $\frac{p_Y}{1-L} < \frac{p_X}{L}$ , or equivalently  $L < \frac{p_X}{p_X + p_Y}$ , we have that  $\underline{v}(L, p_X, p_Y) < \bar{v}(L, p_X, p_Y)$ , hence

non-sensitive data sharing is optimal for intermediate values of  $v_t$ . Moreover,  $\underline{v}(L, p_X, p_Y)$  is increasing in  $L$  and  $\bar{v}(L, p_X, p_Y)$  is decreasing in  $L$  in this case. If  $\frac{p_Y}{1-L} \geq \frac{p_X}{L}$ , then  $\underline{v}(L, p_X, p_Y) = \bar{v}(L, p_X, p_Y) = p_X + p_Y$ . In this case,  $C_t$  either shares full data or no data, depending on how strongly  $C_t$  values privacy. Our first proposition summarizes these insights:

**Proposition 1** *The users' equilibrium data sharing strategy is*

$$S(n, p_X, p_Y; v_t) = \begin{cases} F & \text{if } v_t \leq \underline{v}(L(n), p_X, p_Y), \\ I & \text{if } \underline{v}(L(n), p_X, p_Y) < v_t \leq \bar{v}(L(n), p_X, p_Y), \\ N & \text{if } \bar{v}(L(n), p_X, p_Y) < v_t. \end{cases}$$

If  $\frac{p_Y}{1-L} < \frac{p_X}{L}$ , then  $\underline{v}(L, p_X, p_Y) < \bar{v}(L, p_X, p_Y)$  and  $\frac{\partial \underline{v}}{\partial L} > 0 > \frac{\partial \bar{v}}{\partial L}$ ; otherwise,  $\underline{v}(L, p_X, p_Y) = \bar{v}(L, p_X, p_Y) = p_X + p_Y$ .

Figures 2 to 4 provide graphical illustrations of the users' data sharing decisions when prices are fixed. Figure 2 illustrates that, keeping prices fixed, the privacy valuation cutoff above which a user prefers to share no data falls as the information leakage from non-sensitive data rises, while the privacy valuation cutoff below which full data sharing is optimal rises. At low level of information leakage, the user optimally shares either non-sensitive data or full data, but never no data. At high levels of information leakage, the user optimally shares either no data or full data, but never only non-sensitive data.

Figure 3 illustrates how the data sharing decisions evolve with the state  $n$ , keeping prices fixed. As the firm accumulates more full data vectors, and thus the information leakage from non-sensitive data rises, no data sharing and full data sharing become optimal for wider ranges of privacy valuations. Given that the information leakage from non-sensitive data is bounded above by  $r^2$  (equal to 0.25 in the example shown in the figure), however, non-sensitive data sharing remains optimal for intermediate privacy valuations even in the limit. Finally, Figure 4 shows how the *expected* privacy valuation cutoffs evolve over time keeping prices fixed, illustrating the same qualitative patterns.

In summary, if data prices do not adjust as the information leakage from non-sensitive data rises, then the proportion of digital hermits and the proportion of full data sharers among new users both grow over time, while the proportion of non-sensitive data sharers falls. The intuition




Figure 2: User data sharing strategy as a function of the information leakage from non-sensitive data  $L$  given fixed prices ( $p_X = 1, p_Y = 0.75$ )




Figure 3: User data sharing strategy as a function of  $n$  given fixed prices ( $p_X = 1, p_Y = 0.75, r = 0.5$ )




Figure 4: Expected user data sharing strategy as a function of the period  $t$  given fixed prices ( $p_X = 1, p_Y = 0.75, r = 0.5, V = 6$ )

behind this is as follows. When a user shares full data, they impose a negative externality on future non-sensitive data sharers. The data shared by the former enables the firm to predict unshared sensitive information from non-sensitive data more accurately, which raises the privacy loss suffered by a later user who shares only non-sensitive data. Holding data prices fixed, non-sensitive data sharing therefore becomes less appealing relative to sharing either no data or full data as the firm accumulates more full data vectors.

Next, we assess whether these qualitative insights continue to hold when the firm optimally adjusts data prices as it accumulates more full data vectors and the information leakage from non-sensitive data grows.

### 3.3 Equilibrium data prices and data sharing decisions

#### 3.3.1 Data prices

We first analyze the prices that maximize the firm's current-period payoff (henceforth profit). Given that the current-period profit depends on the state  $n$  only through  $L(n)$ , the myopic profit-maximizing prices will be expressed as functions of  $L$  rather than  $n$  for notational convenience. Using the privacy valuations cutoffs that determine user data sharing decisions, the firm's current-period profit can be written as

$$\begin{aligned} \pi(p_X, p_Y, L) &= (H(\bar{v}(L, p_X, p_Y)) - H(\underline{v}(L, p_X, p_Y)))(1 + L - p_X) \\ &\quad + H(\underline{v}(L, p_X, p_Y))(2 - p_X - p_Y) \\ &= H(\bar{v}(L, p_X, p_Y))(1 + L - p_X) + H(\underline{v}(L, p_X, p_Y))(1 - L - p_Y). \end{aligned}$$

Intuitively, given the level  $L$  of information leakage from non-sensitive data,  $1 + L - p_X$  is the firm's profit from obtaining a user's non-sensitive data and  $1 - L - p_Y$  is the firm's incremental profit from obtaining the user's sensitive data in addition to their non-sensitive data. The former is multiplied by the probability that the user shares non-sensitive data (i.e., they share either only non-sensitive data or full data), the latter is multiplied by the probability that the user shares full data.



**Lemma 2 Myopic profit-maximizing prices.** *The following prices maximize the firm's current-period profit  $\pi(p_X, p_Y, L)$ :*

$$p_X^*(L) = \begin{cases} LV & \text{if } L \leq \frac{1}{2V-1}, \\ \frac{1+L}{2} & \text{if } L > \frac{1}{2V-1}. \end{cases}, p_Y^*(L) = \frac{1-L}{2}.$$

A direct implication of Lemma 2 is that the myopic profit-maximizing price for non-sensitive data is increasing in the information leakage from non-sensitive data, while the myopic profit-maximizing price for sensitive data is decreasing:

$$\frac{\partial p_X^*(L)}{\partial L} > 0 > \frac{\partial p_Y^*(L)}{\partial L}.$$

Greater information leakage from non-sensitive data makes a user's non-sensitive data more valuable to the firm, but it reduces the incremental value of a user's sensitive data. Greater information leakage also means that users require higher compensation to be induced to share non-sensitive data, but less additional compensation to share their sensitive data. It is therefore intuitive that the firm raises the price for non-sensitive data but lowers the price for sensitive data as the information leakage from non-sensitive data rises. Moreover, the myopically optimal total price for a user's full data,  $p_X^*(L) + p_Y^*(L)$ , is weakly increasing in the information leakage from non-sensitive data. It rises at low values of information leakage and is constant at high values of information leakage.

If the firm has discount factor  $\delta = 0$ , the equilibrium prices coincide with the myopic profit-maximizing prices. If the firm is forward-looking ( $\delta > 0$ ), the equilibrium prices solve a dynamic programming problem in which the firm maximizes the expected present discounted value of its current and future profits in every state. Denoting by  $W(n)$  the firm's expected present discounted value of current and future profits in state  $n$ , the value function can be expressed recursively as

$$W(n) = \max_{p_X \geq 0, p_Y \geq 0} (\pi(p_X, p_Y, L(n)) + \delta [H(\underline{v}(L(n), p_X, p_Y)) W(n+1) + (1 - H(\underline{v}(L(n), p_X, p_Y))) W(n)]).$$

Note that the probability of transitioning to the next state depends on the data prices through their effect on the privacy valuation cutoff  $\underline{v}$  below which users share full data.

Our next proposition shows how dynamic considerations affect the firm's equilibrium prices:

**Proposition 2** *In equilibrium, for every  $n \in \mathbb{N}_0$ ,*

(i) *for all  $\delta \geq 0$ , the price for non-sensitive data*

$$\begin{aligned} P_X(n) &= p_X^*(L(n)), \text{ and} \\ P_X(n) &< P_X(n+1). \end{aligned}$$

(ii) *for all  $\delta > 0$ , the price for sensitive data*

$$\begin{aligned} P_Y(n) &< p_Y^*(L(n)), \text{ and} \\ p_Y^*(L(n)) - P_Y(n) &> p_Y^*(L(n+1)) - P_Y(n+1), \end{aligned}$$

*with  $\lim_{n \rightarrow \infty} (p_Y^*(L(n)) - P_Y(n)) = 0$ . Moreover,  $\frac{\partial P_Y(n)}{\partial \delta} < 0$  for all  $P_Y(n) > 0$ .*

(iii) *for all  $\delta > 0$ , the total price for a user's full data*

$$P_X(n) + P_Y(n) < P_X(n+1) + P_Y(n+1).$$

Figures 5 and 6 illustrate these insights. The equilibrium price for non-sensitive data,  $P_X(n)$ , is independent of the discount factor. Intuitively, this is because non-sensitive data on its own does not improve the firm's ability to predict unshared data from non-sensitive data, hence dynamic considerations are absent. The figure also illustrates that the equilibrium price for non-sensitive data rises as the firm accumulates more full data vectors, which follows from our earlier insight that the myopic profit-maximizing price for non-sensitive data increases in the information leakage.

The equilibrium price for sensitive data,  $P_Y(n)$ , lies below the myopically optimal level for all  $\delta > 0$ , because the collection of full data vectors erodes future profits by pushing up the compensation required to induce future users to share non-sensitive data. The firm therefore has an incentive to slow down the collection of full data vectors, which it achieves by offering a lower price for sensitive data. This incentive to slow down learning by distorting the price of sensitive data below the myopically optimal level is stronger when the firm is more patient (higher  $\delta$ ) or has




Figure 5: Equilibrium data prices ( $r = 0.5$ ,  $V = 6$ )




Figure 6: Total price for a user's data in equilibrium ( $r = 0.5$ ,  $V = 6$ )

collected fewer full data vectors so far (lower  $n$ ). The latter is due to our earlier insight that the information leakage from non-sensitive data rises in  $n$  at a decreasing rate (see Lemma 1). It is also worth noting that for high values of the discount factor, the firm may find it optimal to set  $P_Y(0) = 0$  and avoid collecting full data vectors altogether. In this case, no user shares full data in state  $n = 0$ , which implies that  $n = 0$  is an absorbing state.

Figure 6 illustrates that, although the equilibrium price for sensitive data can be decreasing in  $n$ , the total price offered for a user's full data always increases in  $n$ .

### 3.3.2 Data sharing decisions

Recall that in the case of fixed data prices, the privacy valuation cutoff  $\bar{v}$  above which users opt to share no data is decreasing in  $n$ . Users become more inclined to share no data as the firm accumulates more full data vectors. With endogenous prices, it is *a priori* unclear whether this result continues to hold, because the price offered for non-sensitive data rises with  $n$ , which all else equal incentivizes users to share non-sensitive data.

To evaluate the net effect of  $n$  on the equilibrium data sharing choices, we define the privacy valuations cutoffs as functions of the state given the firm's equilibrium data pricing strategy:

$$\bar{v}(n) \equiv \bar{v}(L(n), P_X(n), P_Y(n)),$$

$$\underline{v}(n) \equiv \underline{v}(L(n), P_X(n), P_Y(n)).$$

We obtain the following comparative statics result for  $\bar{v}(n)$ :

**Proposition 3** *In equilibrium, for every  $n \in \mathbb{N}_0$ ,*

$$\bar{v}(n) \geq \bar{v}(n+1),$$

*with a strict inequality if and only if  $L(n+1) > \frac{1}{2V-1}$ .*

Hence, as in the case of fixed prices, users become more inclined to share no data as the firm accumulates more full data vectors and the information leakage from non-sensitive data rises.

Turning to the equilibrium privacy valuation cutoff  $\underline{v}(n)$  below which users share all data, we can first note that  $\underline{v}(n+1) = \underline{v}(n) = \frac{1}{2}$  for all  $n$  if the firm is myopic ( $\delta = 0$ ). Moreover, extensive numerical simulations<sup>2</sup> show that for  $\delta > 0$ , the equilibrium threshold below which users share full data rises as the information leakage from non-sensitive data rises:

$$\underline{v}(n) < \underline{v}(n+1) \text{ for all } n \in \mathbb{N}_0 \text{ if } \delta > 0.$$

Intuitively, this is because, while the price for sensitive data always decreases in  $n$  when  $\delta = 0$ , it either increases or decreases less rapidly in  $n$  when  $\delta > 0$ . Compared to the myopic-firm case, users therefore become relatively more inclined to share full data as  $n$  rises.

Figure 7 illustrates the equilibrium data sharing decisions as a function of  $n$  for different values of the discount factor. As in the case of fixed prices, the proportion of new users who share no data and the proportion of new users who share full data in the population rise as the firm accumulates more full data vectors and thus the information leakage from non-sensitive data grows.

Figure 8 illustrates the same qualitative patterns in how the expected equilibrium data sharing privacy valuation cutoffs evolve over time. The expected privacy valuation cutoff above which users share no data falls over time, and the expected privacy valuation cutoff below which users share no data rises. Moreover, for a given  $t$ , the expected  $\bar{v}(n)$  is higher when the firm is more patient. The reason is that, as discussed earlier, the firm's incentive to slow down learning is stronger, and hence the information leakage from non-sensitive data rises less fast over time, when the firm is more patient.

---

<sup>2</sup>Specifically, we simulated the equilibrium for each point on the following grid of parametrizations:  $\delta \in (0, 0.01, \dots, 0.99)$ ,  $V \in (2, 4, \dots, 20)$ ,  $r \in (0.05, 0.10, \dots, 0.95)$




Figure 7: Equilibrium user data sharing decisions ( $r = 0.5$ ,  $V = 6$ )

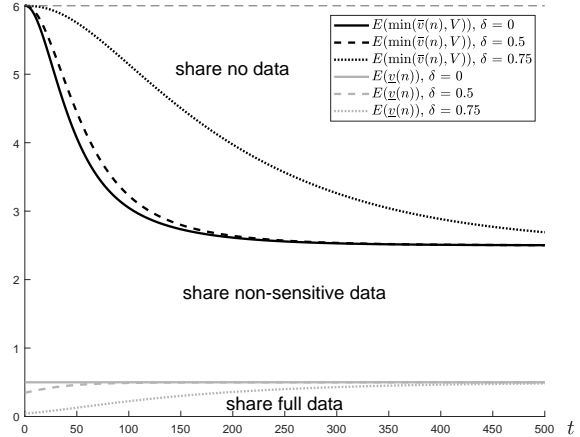


Figure 8: Expected equilibrium user data sharing decisions over time ( $r = 0.5$ ,  $V = 6$ )

## 4 Conclusion

Data about some aspects of an individual can reveal other, more privacy sensitive, aspects of that same individual. As firms collect more data and learn about the correlations between different dimensions of individual data, the possibility that they can use seemingly innocuous data to predict information that people prefer to keep private grows. Our paper shows that this leads to a polarization of users' data sharing decisions. As the firm's ability to infer sensitive data from non-sensitive data improves, more privacy-protective people will decide to avoid sharing any data with firms, even data that on the face of it is not sensitive, while less privacy-protective people, realizing that there is little point keeping only some data private, will share all their data with firms. This may disincentivize firms from collecting sensitive data in the first place, and they end up collecting it more slowly by offering a lower price.

There are of course limitations to this research. First, we do not model the root cause of why it is that people are more uncomfortable about some types of data than others. Second, we assume that consumers are aware of the firm's ability to infer sensitive information from non-sensitive data. Third, our treatment of how data markets work is intentionally simple. For example, we focus on short-lived users. Notwithstanding these limitations, however, we feel this paper makes a useful first step in understanding how the correlations between data and the inferences that can be drawn about them can lead people to withdraw from the digital world.

## References

- Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2022). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics* 14(4), 218–56.
- Acquisti, A., C. Taylor, and L. Wagman (2016). The economics of privacy. *Journal of Economic Literature* 54(2), 442–92.
- Acquisti, A. and H. R. Varian (2005). Conditioning prices on purchase history. *Marketing Science* 24(3), 367–381.
- Bergemann, D., A. Bonatti, and T. Gan (2022). The economics of social data. *The RAND Journal of Economics* 53(2), 263–296.
- Choi, J. P., D.-S. Jeon, and B.-C. Kim (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics* 173, 113–124.
- Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics* 24(3), 461–483.
- Grossman, S. J. and O. D. Hart (1980). Disclosure laws and takeover bids. *The Journal of Finance* 35(2), 323–334.
- Ichihashi, S. (2020). Dynamic privacy choices. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 539–540.
- Ichihashi, S. (2021). The economics of data externalities. *Journal of Economic Theory* 196(C).
- Kosinski, M., D. Stillwell, and T. Graepel (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences* 110(15), 5802–5805.
- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics* 12(2), 380–391.
- Prince, J. T. and S. Wallsten (2022). How much is privacy worth around the world and across platforms? *Journal of Economics & Management Strategy* 31(4), 841–861.

Taylor, C. R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics* 35(4), 631–650.

## Appendix: Proofs

**Proof of Lemma 1.** The information leakage from non-sensitive data is

$$L(n) = \sum_{k=0}^n P(k|n) \widehat{\rho}(k|n)^2,$$

where  $P(k|n)$  is the prior probability that  $k$  out of  $n$  full data vectors indicate positive correlation (i.e.,  $(x, y) \in \{(0, 0), (2\sigma_X, 2\sigma_Y)\}$ ):<sup>3</sup>

$$P(k|n) = \frac{1}{2} \binom{n}{k} \left( \left( \frac{1+r}{2} \right)^k \left( \frac{1-r}{2} \right)^{n-k} + \left( \frac{1-r}{2} \right)^k \left( \frac{1+r}{2} \right)^{n-k} \right),$$

and  $\widehat{\rho}(k|n)$  is the firm's estimate of  $\rho$  given  $k$  and  $n$ :

$$\widehat{\rho}(k|n) = E[\rho|k, n] = r \frac{\left(\frac{1+r}{2}\right)^k \left(\frac{1-r}{2}\right)^{n-k} - \left(\frac{1-r}{2}\right)^k \left(\frac{1+r}{2}\right)^{n-k}}{\left(\frac{1+r}{2}\right)^k \left(\frac{1-r}{2}\right)^{n-k} + \left(\frac{1-r}{2}\right)^k \left(\frac{1+r}{2}\right)^{n-k}}.$$

Using  $\widehat{\rho}(k|n) = \frac{1+\widehat{\rho}(k|n)}{2} \widehat{\rho}(k+1|n+1) + \frac{1-\widehat{\rho}(k|n)}{2} \widehat{\rho}(k|n+1)$  and  $P(k|n+1) = \frac{1+\widehat{\rho}(k-1|n)}{2} P(k-1|n) + \frac{1-\widehat{\rho}(k|n)}{2} P(k|n)$ , we obtain

$$\begin{aligned} L(n) &= \sum_{k=0}^n P(k|n) \left( \frac{1+\widehat{\rho}(k|n)}{2} \widehat{\rho}(k+1|n+1) + \frac{1-\widehat{\rho}(k|n)}{2} \widehat{\rho}(k|n+1) \right)^2 \\ L(n+1) &= \sum_{k=0}^n P(k|n) \left( \frac{1+\widehat{\rho}(k|n)}{2} \widehat{\rho}(k+1|n+1)^2 + \frac{1-\widehat{\rho}(k|n)}{2} \widehat{\rho}(k|n+1)^2 \right). \end{aligned}$$

Since  $\frac{1+\widehat{\rho}(k|n)}{2} = 1 - \frac{1-\widehat{\rho}(k|n)}{2} \in (0, 1)$ ,  $\widehat{\rho}(k+1|n+1) \neq \widehat{\rho}(k|n+1)$ , and the function  $x^2$  is strictly convex,

$$\frac{1+\widehat{\rho}(k|n)}{2} \widehat{\rho}(k+1|n+1)^2 + \frac{1-\widehat{\rho}(k|n)}{2} \widehat{\rho}(k|n+1)^2 > \left( \frac{1+\widehat{\rho}(k|n)}{2} \widehat{\rho}(k+1|n+1) + \frac{1-\widehat{\rho}(k|n)}{2} \widehat{\rho}(k|n+1) \right)^2$$

by Jensen's inequality. Hence,  $L(n+1) > L(n)$ .

---

<sup>3</sup>For  $k > n$ ,  $P(k|n) = 0$ .



Define the ‘Jensen gap’ by

$$J(k, n) = \frac{1 + \widehat{\rho}(k|n)}{2} \widehat{\rho}(k+1|n+1)^2 + \frac{1 - \widehat{\rho}(k|n)}{2} \widehat{\rho}(k|n+1)^2 - \left( \frac{1 + \widehat{\rho}(k|n)}{2} \widehat{\rho}(k+1|n+1) + \frac{1 - \widehat{\rho}(k|n)}{2} \widehat{\rho}(k|n+1) \right)^2,$$

which can be simplified to

$$J(k, n) = \frac{16r^4}{\left( \left( \frac{1+r}{1-r} \right)^{\frac{n}{2}-k} + \left( \frac{1-r}{1+r} \right)^{\frac{n}{2}-k} \right)^4 - r^2 \left( \left( \frac{1+r}{1+r} \right)^{n-2k} - \left( \frac{1-r}{1+r} \right)^{n-2k} \right)^2}.$$

We then have that

$$\begin{aligned} L(n+1) - L(n) &= \sum_{k=0}^n P(k|n) J(k, n) = \sum_{k=0}^{n+1} P(k|n) J(k, n), \\ L(n+2) - L(n+1) &= \sum_{k=0}^{n+1} P(k|n+1) J(k, n+1). \end{aligned}$$

Straightforward calculations show that  $J(k+1, n) < J(k, n)$  and  $J(k, n+1) < J(k, n)$ . Hence,

$$L(n+1) - L(n) - [L(n+2) - L(n+1)] > \sum_{k=0}^{n+1} P(k|n) J(k, n+1) - \sum_{k=0}^{n+1} P(k|n+1) J(k, n+1). \quad (1)$$

The distribution  $P(k|n+1)$  first-order stochastically dominates  $P(k|n)$ , because  $\sum_{m=0}^k P(k|n) > \sum_{m=0}^k P(k|n+1)$  for all  $k \in \{1, 2, \dots, n\}$ , both sums are zero for  $k = 0$ , and both sums are equal to 1 for  $k \geq n+1$ . First-order stochastic dominance together with  $J(k+1, n) < J(k, n)$  imply that the right-hand side of (1) is strictly positive, and thus  $L(n+1) - L(n) > L(n+2) - L(n+1)$ .

For  $n = 0$ ,  $L(0) = P(0|0) \widehat{\rho}(0|0)^2 = E[\rho]^2 = 0$ . Moreover, by the law of large numbers,

$$\lim_{n \rightarrow \infty} L(n) = \frac{1}{2} (-r)^2 + \frac{1}{2} (r)^2 = r^2.$$

■

**Proof of Lemma 2.** First, consider  $L = 0$ . In this case,

$$\pi(p_Y, p_X, L) = \min \left\{ \frac{p_Y}{V}, 1 \right\} (1 - p_Y) + 1 - p_X,$$

which is maximized for  $(p_X, p_Y) = (0, \frac{1}{2})$ .

For the remainder of the proof, consider  $0 < L < 1$ . First, suppose that  $\frac{p_Y}{1-L} < \frac{p_X}{L}$ , and thus  $\underline{v} < \bar{v}$ , at the solution of the firm's profit maximization problem. The firm's problem then is

$$\max_{p_X \geq 0, p_Y \geq 0} \left( (1-L) H \left( \frac{p_Y}{1-L} \right) \left( 1 - \frac{p_Y}{1-L} \right) + LH \left( \frac{p_X}{L} \right) \left( \frac{1+L}{L} - \frac{p_X}{L} \right) \right).$$

If  $L < \frac{1}{2V-1}$ , the profit function is increasing in  $p_X$  for all  $p_X \leq VL$  and decreasing in  $p_X$  for  $p_X > VL$ , hence profit is maximized at  $p_X = VL$ . If  $L \geq \frac{1}{2V-1}$ , profit is maximized at  $p_X = \frac{1+L}{2} < V$ . Moreover, for all  $L \in (0, 1)$ , it is optimal to set  $p_Y = \frac{1-L}{2}$ . In both cases,  $\frac{p_Y}{1-L} < \frac{p_X}{L}$  at the solution.

It remains to show that the firm cannot do better by setting prices such that  $\frac{p_Y}{1-L} \geq \frac{p_X}{L}$ , and thus  $\underline{v} = \bar{v}$ . At such prices, the firm's current-period profit would be  $H(p_X + p_Y) (2 - (p_X + p_Y))$ . By revealed preferences, the firm can earn more by setting the profit-maximizing prices derived earlier:

$$\begin{aligned} & \max_{p_X \geq 0, p_Y \geq 0} \left( (1-L) H \left( \frac{p_Y}{1-L} \right) \left( 1 - \frac{p_Y}{1-L} \right) + LH \left( \frac{p_X}{L} \right) \left( \frac{1+L}{L} - \frac{p_X}{L} \right) \right) \\ & \geq \max_{\hat{p}} \left( (1-L) (H(\hat{p}) (1 - \hat{p})) + LH(\hat{p}) \left( \frac{1+L}{L} - \hat{p} \right) \right) \\ & = \max_{\hat{p}} (H(\hat{p}) (2 - \hat{p})) \geq H(p_X + p_Y) (2 - (p_X + p_Y)). \end{aligned}$$

■

**Proof of Proposition 2.** The value function can be written recursively as

$$(1 - \delta) W(n) = \max_{p_X \geq 0, p_Y \geq 0} (\pi(p_X, p_Y, L(n)) + \delta (H(\underline{v}(L(n), p_X, p_Y)) [W(n+1) - W(n)])).$$

Suppose that  $\frac{p_Y(n)}{1-L(n)} < \frac{p_X(n)}{L(n)} \leq V$  for all  $n$ , so that  $\underline{v}(L(n), p_X, p_Y) = \frac{p_Y}{V(1-L)}$ . We will show at the end of the proof that this is indeed the case. Since  $\pi(p_X, p_Y, L(n))$  is separable in  $p_X$  and  $p_Y$ , it then follows that  $p_X(n) = p_X^*(L(n))$  and

$$(1 - \delta) W(n) = \max_{p_Y \geq 0} \left( \pi(p_X^*(n), p_Y, L(n)) + \delta \left( \frac{p_Y}{V(1-L(n))} [W(n+1) - W(n)] \right) \right). \quad (2)$$

Any interior solution  $P_Y(n) > 0$  must satisfy the first-order condition

$$P_Y(n) = \frac{1 - L(n)}{2} - \frac{\delta}{2} [W(n) - W(n+1)] = p_Y^*(L(n)) - \frac{\delta}{2} [W(n) - W(n+1)]. \quad (3)$$

Moreover, Lemmas 1 and 2 imply that  $P_X(n+1) > P_X(n)$ .

We now show that  $W(n+1) - W(n) > W(n+2) - W(n+1) > 0$  and  $\lim_{n \rightarrow \infty} (W(n+1) - W(n)) = 0$ . First, note that the firm's problem in (2) can be reformulated as one in which the firm sets  $\underline{v}$ 's instead of  $p_Y$ 's:

$$(1 - \delta) W(n) = \max_{\underline{v} \geq 0} (\hat{\pi}(\underline{v}, n) + \delta (H(\underline{v}) [W(n+1) - W(n)])), \quad (4)$$

where

$$\hat{\pi}(\underline{v}, n) = \pi(p_X^*(L(n)), \underline{v}(1 - L(n)), L(n)) = (1 - L(n)) \frac{\underline{v}(1 - \underline{v})}{V} + \frac{1}{V} \frac{(1 + L(n))^2}{4L(n)}.$$

Since this function is strictly decreasing and strictly convex in  $L(n)$  for any  $\underline{v}$ , Lemma 1 implies that for any  $\underline{v}$ ,

$$\hat{\pi}(\underline{v}, n) - \hat{\pi}(\underline{v}, n+1) > \hat{\pi}(\underline{v}, n+1) - \hat{\pi}(\underline{v}, n+2) > 0. \quad (5)$$

Next, denote by  $\hat{\mathbf{v}}_n = \left\{ \frac{P_Y(n)}{1-L(n)}, \frac{P_Y(n+1)}{1-L(n+1)}, \dots \right\}$  the equilibrium sequence of  $\underline{v}$ 's starting in state  $n$ . Denote by  $W(n | \hat{\mathbf{v}}_m)$  the expected present discounted value of profits if the state is  $n$  and the firm uses the sequence  $\hat{\mathbf{v}}_m$  from the current state onwards, where  $m$  can be different from  $n$ . By definition,  $W(n) = W(n | \hat{\mathbf{v}}_n)$ . Note that the transition probabilities in  $W(n | \hat{\mathbf{v}}_m)$  depend on  $m$ , but not on  $n$ . Hence, for any  $\Delta \in \mathbb{N}_0^+$ , the probability weight assigned to the profit  $\hat{\pi}(\underline{v}(m + \Delta), n + \Delta)$  in  $W(n | \hat{\mathbf{v}}_m)$  is the same as the probability weight assigned to the profit  $\hat{\pi}(\underline{v}(m + \Delta), n' + \Delta)$  in  $W(n' | \hat{\mathbf{v}}_m)$ .

The inequalities in (5) then imply that

$$W(n | \hat{\mathbf{v}}_{n+1}) > W(n+1 | \hat{\mathbf{v}}_{n+1}), \text{ and} \quad (6)$$

$$2W(n+1 | \hat{\mathbf{v}}_{n+1}) < W(n | \hat{\mathbf{v}}_{n+1}) + W(n+2 | \hat{\mathbf{v}}_{n+1}). \quad (7)$$

Since  $W(n) = W(n|\widehat{\mathbf{v}}_n) \geq W(n|\widehat{\mathbf{v}}_m)$  for all  $n$  and  $m$ , (6) implies that

$$W(n) > W(n+1), \quad (8)$$

and (7) implies that

$$\begin{aligned} 2W(n+1) &< W(n|\widehat{\mathbf{v}}_{n+1}) + W(n+2|\widehat{\mathbf{v}}_{n+1}) \\ &< W(n|\widehat{\mathbf{v}}_n) + W(n+2|\widehat{\mathbf{v}}_{n+2}) = W(n) + W(n+2). \end{aligned} \quad (9)$$

Finally,  $\lim_{n \rightarrow \infty} (\widehat{\pi}(\underline{v}, n) - \widehat{\pi}(\underline{v}, n+1)) = 0$  for all  $\underline{v}$  by Lemma 1, which implies that  $\lim_{n \rightarrow \infty} (W(n|\widehat{\mathbf{v}}_n) - W(n+1)) = 0$ . Given  $W(n+1|\widehat{\mathbf{v}}_n) \leq W(n+1)$  and (8), it follows that

$$\lim_{n \rightarrow \infty} (W(n) - W(n+1)) = 0. \quad (10)$$

Results (8), (9), and (10) together with the first-order condition in (3) imply that

$$\begin{aligned} P_Y(n) &< p_Y^*(L(n)), \\ p_Y^*(L(n)) - P_Y(n) &> p_Y^*(L(n+1)) - P_Y(n+1), \text{ and} \\ \lim_{n \rightarrow \infty} (p_Y^*(L(n)) - P_Y(n)) &= 0. \end{aligned}$$

To complete the proof of part (ii), we apply the implicit function theorem to (3) and use the envelope theorem to obtain

$$\begin{aligned} \frac{\partial P_Y(n)}{\partial \delta} &= - \left( \frac{1}{2} [W(n) - W(n+1)] + \frac{\delta}{2} \frac{\partial [W(n) - W(n+1)]}{\partial \delta} \right) \\ &= - \frac{1}{2} \left( 1 + \frac{\delta}{1-\delta} (1 - H(\underline{v}(n))) \right) [W(n) - W(n+1)] \\ &\quad - \frac{\delta}{2(1-\delta)} H(\underline{v}(n+1)) [W(n+1) - W(n+2)] \\ &< 0. \end{aligned}$$

For part (iii) of the proposition, note that the total equilibrium price

$$P_X(n) + P_Y(n) = p_X^*(L(n)) + p_Y^*(L(n)) - \frac{\delta}{2} [W(n) - W(n+1)].$$

Lemmas 1 and 2 together with (9) imply that this expression is strictly increasing in  $n$ .

It remains to show that the seller cannot do better by setting prices such that  $\frac{P_Y(n)}{1-L(n)} \geq \frac{P_X(n)}{L(n)}$  for some or all  $n$ . Suppose (in negation) that the equilibrium prices satisfy  $\frac{P_Y(n)}{1-L(n)} \geq \frac{P_X(n)}{L(n)}$  in state  $n$  and denote the total price in that state by  $P_T(n) = P_X(n) + P_Y(n)$ . The seller then earns a current-period profit of  $H(P_T(n))(2 - P_T(n))$  and the probability of transitioning to the next state is  $H(P_T(n))$ . Now suppose the seller instead sets the prices  $(\tilde{p}_X, \tilde{p}_Y)$  so that  $\tilde{p}_Y = (1 - L(n)) P_T(n)$  and  $\tilde{p}_X$  is set to maximize current-period profit subject to the constraint  $\frac{\tilde{p}_X}{L(n)} \geq \frac{\tilde{p}_Y}{1-L(n)} = P_T(n)$ . The probability of transitioning to the next state then remains  $H(P_T(n))$ , but the seller's current-period profit weakly increases because

$$\begin{aligned} & \max_{\tilde{p}_X \geq P_T(n)} \left( (1 - L(n)) H(P_T(n)) (1 - P_T(n)) + L(n) H\left(\frac{\tilde{p}_X}{L(n)}\right) \left(\frac{1 + L(n)}{L(n)} - \frac{\tilde{p}_X}{L(n)}\right) \right) \\ & \geq (1 - L(n)) H(P_T(n)) (1 - P_T(n)) + L(n) H(P_T(n)) \left(\frac{1 + L(n)}{L(n)} - P_T(n)\right) \\ & = H(P_T(n)) (2 - P_T(n)). \end{aligned}$$

■

**Proof of Proposition 3.** By Proposition 2,  $P_X(n) = p_X^*(L(n))$  and  $P_Y(n) \leq p_Y^*(L(n))$ , which implies  $\underline{v}(n) < \bar{v}(n)$ . Hence, using Lemma 2,

$$\bar{v}(n) = \frac{p_X^*(L(n))}{L(n)} = \begin{cases} V & \text{if } L(n) \leq \frac{1}{2V-1} \\ \frac{1+L}{2L} & \text{if } L(n) > \frac{1}{2V-1} \end{cases}.$$

It follows from the expression for  $\bar{v}(n)$  that  $\bar{v}(n+1) = \bar{v}(n)$  if  $L(n+1) \leq \frac{1}{2V-1}$  and  $\bar{v}(n+1) < \bar{v}(n)$  if  $L(n+1) > \frac{1}{2V-1}$ . ■