The Incentives to Preserve Digital Materials:

Roles, Scenarios, and Economic Decision-Making

Brian F. Lavoie
(lavoie@oclc.org)

Office of Research
(http://www.oclc.org/research/)

OCLC Online Computer Library Center, Inc.
(http://www.oclc.org/)

April 2003

Economic issues are a principal component of the research agenda for digital preservation. Economics is fundamentally about incentives, so a study of the economics of digital preservation should begin with an examination of the incentives to preserve. Securing the long-term viability and accessibility of digital materials requires an appropriate allocation of incentives among key decision-makers in the digital preservation process. But the circumstances under which digital preservation takes place often lead to a misalignment of preservation objectives and incentives. Identifying circumstances where insufficient incentives to preserve are likely to prevail, and how this can be remedied, are necessary first steps in developing economically sustainable digital preservation activities.

http://www.oclc.org/research/projects/digipres/incentives-dp.pdf

Suggested citation:
Lavoie, Brian F. (2003) *The Incentives to Preserve Digital Materials: Roles, Scenarios, and Economic Decision-Making*. White paper published electronically by OCLC Research. Available online at:
http://www.oclc.org/research/projects/digipres/incentives-dp.pdf

ACKNOWLEDGEMENTS

The author has benefited from the insightful comments of a number of invited reviewers, including Meg Bellinger, Lorcan Dempsey, Clive Field, Dale Flecker, Clifford Lynch, Robert Martin, Ann Okerson, Edward O'Neill, Helen Shenton, Abby Smith, Hal Varian, and Deborah Woodyard. Special thanks go to Paul Conway and Donald Waters, who provided particularly detailed and thorough comments. The contributions of all reviewers have improved the paper significantly. Any remaining errors of omission or commission are, of course, the responsibility of the author.

EXECUTIVE SUMMARY

Digital preservation tops the agendas of many institutions in the cultural heritage and information management communities. As the number of digital assets which they own outright, or for which they are at least stakeholders, continues to expand, the need to take measures to secure their long-term availability grows commensurately. Technical challenges remain in regard to achieving this objective. But as digital preservation moves beyond the realm of small-scale, experimental projects to become a routine component of a digital asset's life-cycle management, the question of how it can be shaped into an economically sustainable process begins to overshadow other concerns.

The fundamental economic issue associated with digital preservation concerns the *incentives* to preserve digital materials. The incentives to preserve can be characterized as perceived motivation sufficient to 1) induce a party to recognize a need to take action to secure the long-term viability of digital materials in which they are a stakeholder, and 2) induce a party to develop and implement technologies aimed at ensuring the long-term viability of digital materials. These incentives impact three key economic decision-makers in the digital preservation process:

- **Rights Holder:** holds the intellectual property rights to the digital materials
- **Archive:** provides services to ensure the long-term preservation and accessibility of digital materials
- **Beneficiary:** benefits from the long-term retention of digital materials

It is important to note that the above are *roles*, rather than distinct entities. Moreover, a single entity can subsume one, two, or all of these roles. Given these three decision-making roles, it is possible to define five organizational models, or scenarios, under which digital preservation activities might take place:

- **Centripetal Model:** Rights Holder, Archive and Beneficiary are the same entity
- **Centrifugal Model:** Rights Holder, Archive, and Beneficiary are all separate entities
- **Supply-Side Model:** Rights Holder and Archive are the same entity; Beneficiary is separate
- **Demand-Side Model:** Rights Holder and Beneficiary are the same entity; Archive is separate
- **Consolidated Model:** Archive and Beneficiary are the same entity; Rights Holder is separate

The nature and extent of the incentives to preserve are impacted heavily by the structure of the particular organizational model under which digital preservation takes place. This is best illustrated by considering the role of incentives in the formation of an aftermarket in digital preservation services. Aftermarkets provide goods and services that enhance the value and longevity of a durable good – such as digital information – purchased in the primary market. Given the fragility of digital resources, an aftermarket

in digital preservation services would seem likely to emerge. Yet such a market has been slow to develop. It is useful to consider whether, in light of the five organizational models, sufficient incentives exist on the part of relevant decision-makers to support the rapid development of an aftermarket in digital preservation services.

Digital preservation embodies at least three characteristics which could potentially diminish the incentives for decision-makers to take the steps necessary to secure the long-term retention of digital materials. The first characteristic is the fact that those who own digital materials are, in many cases, different from those who would benefit from their preservation: in other words, the Rights Holder is an entity distinct from the Beneficiary. Since the party responsible for carrying out digital preservation does not capture the associated benefits, it is unlikely they will take these benefits into account in the course of their decision-making process, resulting in a less-than-optimal amount of preservation activity taking place. Scenarios where the Rights Holder is a separate entity from the Beneficiary include the Centrifugal, Supply-Side, and Consolidated Models.

Circumstances where the digital materials to be preserved are non-unique can also diminish the incentives to preserve. If multiple institutions own copies of the same digital resource, there is little incentive to be the first to take action to preserve the resource; rather, each institution has the incentive to wait for one of the others to act. If one institution goes to the trouble and expense of preserving the resource, the others can "free-ride" on the benefits – once the resource is preserved for one institution, it is perceived as being preserved for all, since the preserved copy is substitutable for all other copies. In the context of the organizational models, situations where the Rights Holder and the Beneficiary are the same entity, and there are multiple institutions filling these dual roles, are most prone to the free-riding problem. These conditions can occur in the Centripetal Model and the Demand-Side Model.

Finally, the incentives to preserve can be impacted by the presence of heterogeneity in the demand for digital preservation services. Division of the market into those who demand a "high-end" digital preservation service (e.g., preservation in perpetuity, or preservation of the original "look and feel" of the digital resource) and those who demand a "low-end" service (preservation for a relatively short time-horizon, or preservation of the resource's intellectual content only) creates difficulties for suppliers of digital preservation services in several ways. First, segmentation in demand can serve to limit the size of the market for a particular level of preservation service, or even for preservation services in general. This in turn could prevent the Archive from realizing economies of scale in the provision of one or more levels of service, with adverse consequences for the costs, and ultimately, the returns, associated with the Archive's activities. Also, the inability to exclude customers for the low-end service from enjoying "spillover" benefits from customers of the "high-end" service could reduce the potential revenues recoverable from the market, and thereby diminish the incentive for the Archive to supply digital preservation services. These circumstances are most likely to occur in situations where the Archive is an entity distinct from the Beneficiary, there are multiple Beneficiaries, and there is a varying perception among these Beneficiaries as

to the level of preservation service needed. These circumstances can occur in the Centrifugal, Supply-Side, and Demand-Side Models.

In scenarios where incentives prove insufficient to motivate decision-makers to carry out digital preservation, there is scope for adopting corrective measures to enhance these incentives. In the case where Rights Holders are entities distinct from Beneficiaries, one option is to subsidize the Rights Holder to encourage them to carry out digital preservation; the exertion of bargaining power on the part of Beneficiaries in negotiating with the Rights Holder is another strategy. The problem of free-riding can be avoided by devising means to exclude non-payers from the benefits of preservation, such as making preservation and access inseparable. The problem of heterogeneity in demand among Beneficiaries of digital preservation services can be remedied by offering differing levels of preservation services – but only if 1) doing so does not inhibit the realization of economies of scale by limiting the size of the market, and 2) the market can be segmented in such a way that customers for low-end services are excluded from realizing "spillover benefits" from customers for high-end services.

A thorough understanding of the nature and extent of the incentives to preserve is fundamental to more focused research into the economics of digital preservation. Future research activity should include the accumulation and synthesis of digital preservation case studies; the development of appropriate policies for enhancing incentives based on the characteristics of the underlying organizational model; characterizing and analyzing the structure of aftermarkets associated with digital preservation services; and devising sustainable pricing strategies for digital preservation services.

I. ECONOMICS, INCENTIVES, AND PRESERVATION

The Civil War has been a subject of enduring interest to the American public, and is arguably the most studied event in American history. One factor motivating this keen interest is that the Civil War is one of only three conflicts waged on American soil[1], and by far the largest in scale and scope. Given the physical proximity of its events, the "primary sources" constituting the historical record of the Civil War are readily at hand, extending beyond documents and physical artifacts to include the battlefields themselves. Vast national parks, such as those at Gettysburg and Vicksburg, are testimony to the national commitment to preserve the locations where the War's most important events transpired.

But preservation of Civil War battlefields is not always easy to achieve. In 1864-65, a relatively obscure campaign was waged on Virginia's Bermuda Hundred peninsula. The campaign, little more than a footnote in the military history of the Civil War, nevertheless left a highly visible impact on the land over which it was waged, in the form of entrenchments, earthworks, and fortifications. Many of these survive to this day. Some believe these remnants of the Bermuda Hundred campaign to be important historical artifacts. One difficulty emerged, however, in the course of securing their preservation: most of the sites were on private property.[2]

It therefore fell to advocates of preserving the Bermuda Hundred battlefield sites to persuade property owners (often land developers) to participate in the preservation process. But this proved to be a challenging task. "There are no laws protecting the earthworks," notes one account. "Property owners can level them … But the earthworks in the Bermuda Hundred are on no one's list of most endangered Civil War sites. Too many other nearby sites are more important, or better known, or more visited, or connected to a more glamorous campaign."[3] Many of the sites have already been destroyed, and those that remain "have a slim chance for survival".[4]

Not all of the rescue efforts were unsuccessful, however. One preservationist, who has intervened to save a number of the Bermuda Hundred sites, observes that "… for the sites I've saved, it has been a win-win situation for everybody. The sites are preserved, and the developers get tax write-offs and good press."[5]

The Bermuda Hundred battlefield sites illustrate a critical element in any preservation effort: preservation objectives must be aligned with the incentives for relevant decision-makers to carry them out. The lack of a compelling interest on the part of property owners to act to preserve the portions of the entrenchments and earthworks crossing their land often proved fatal to the goal of historical preservation, despite the fact that viable cost and revenue models were available (the costs associated with

[1] Since the Declaration of Independence in 1776.
[2] The story of the efforts to preserve the Bermuda Hundred battlefield sites is described in Zeller, B. (2002) "Where the Earthworks Meet the Road" *Civil War Times Illustrated* Vol. XLI No. 1, p.18, 58-67
[3] Zeller, p. 58
[4] Ibid
[5] Ibid

preserving historical sites are well understood; revenue sources, such as admission fees or invited donations, have been widely employed in similar circumstances); the cost itself was not necessarily prohibitive; an identifiable constituency would benefit (e.g., Civil War historians and enthusiasts), and the necessary expertise existed to carry out the preservation process.

The cases where preservation was successfully carried out are exceptions that prove the rule – i.e., the primacy of incentives. In these instances, as in all the others, incentives to preserve either did not exist, or were overshadowed by more powerful incentives associated with other uses for the historical sites. But the reticence on the part of key decision-makers to undertake preservation was overcome by supplying additional inducements – a reduction in their tax burden and the opportunity to enhance their public image – that served to enhance the incentives to preserve *vis-à-vis* other choices (like doing nothing, or putting the land to some other use), and ultimately, secured participation in the preservation process.

Economic conditions often are such that decision-makers find little motivation to address preservation issues. This is a theme that will be re-visited throughout this paper. But the instances where Bermuda Hundred sites were successfully preserved illustrate another theme: that measures can be adopted to bring preservation incentives in line with preservation objectives. These measures can assume a variety of forms. Intervention by a third-party funder is clearly one approach – tax relief (as in the case of the Bermuda Hundred sites), subsidies, or philanthropic grants all fall into this category. But other strategies are available as well: in particular, negotiation, persuasion, or simply raising awareness among stakeholders in the preservation process may also enhance the incentives to preserve. Non-economic incentives, such as an improved public image, the cultivation of goodwill in the community, or even mutual recognition that, without action, a part of the cultural or scholarly record will be lost, can be powerful motivators.

The importance of incentives is a principle that extends to preservation in any form. Digital preservation – the art and science of perpetuating information in digital form over the long-term – is no exception. Discourse on the challenges posed by digital preservation often identifies economics as one of the principal obstacles to establishing sustainable digital preservation activities. But economics, and by extension, the economics of digital preservation is, fundamentally, about incentives. As one economist observes: "Most of economics can be summarized in four words: 'People respond to incentives.' The rest is commentary."[6]

* * *

The long-term preservation of the intellectual and cultural record of society has occupied librarians, archivists, and museum curators for centuries. But this issue has recently been re-invigorated by the proliferation of technologies for creating and storing information in digital form. There is general consensus that digital preservation represents both an imperative and a challenge: an imperative, because of the rapidly

[6] Landsburg, S. (1993) *The Armchair Economist* (The Free Press, New York) p. 3

expanding body of information in digital form, and a challenge, because the process of preserving that information over the long term, as well as maintaining its accessibility, is not yet well understood.

The immediacy of this problem – i.e., the fact that many libraries and other cultural heritage institutions are already managing (and expanding) sizable digital collections – has motivated an ambitious, community-wide research agenda aimed at an improved understanding of what digital preservation entails. Great strides have been achieved in many areas, including digital repository architectures, preservation metadata, preservation strategies, the longevity of digital storage media, and risk assessment of digital formats.[7]

It is widely acknowledged that some of the most vexing issues associated with digital preservation are economic. Yet a systematic study of the economics of digital preservation has thus far failed to materialize, yielding pride of place to the *technical* challenges posed by the long-term retention of digital materials. A recent NSF-sponsored workshop, involving participants from across the digital preservation community, concluded that "[m]uch digital preservation research has focused on technical problems and technological solutions without careful analysis of the social, organizational, and economic mechanisms that have to be in place to make preservation possible and sustainable." The workshop designated economic and policy models as one of four key research areas in digital preservation.[8]

The economic issues associated with digital preservation can be summarized by the following question: will the appropriate decision-makers allocate the resources necessary to ensure the long-term viability and accessibility of digital materials?[9] From one perspective, the issue can be expressed in terms of the *nature and magnitude* of the resources required to carry out digital preservation – in other words, enumeration, and where possible, quantification, of the costs involved in retaining digital materials over the long term. A related issue concerns the development of strategies for re-couping these costs through the generation of revenues. Viewed from this perspective, the goal of research into the economics of digital preservation is to develop cost models, and to a lesser extent, business models, that can be used as tools to evaluate whether digital

---

[7] See, for example, the Open Archival Information System (OAIS) reference model (http://ssdoo.gsfc.nasa.gov/nost/isoas/us/overview.html), the OCLC/RLG Working Group on Preservation Metadata (http://www.oclc.org/research/pmwg/), the CAMiLEON project (http://www.si.umich.edu/CAMILEON/), the resources listed under the "Physical Format Digital Media" section of the PADI Web site (http://www.nla.gov.au/padi/topics/52.html), and Lawrence, et al. (2000) "Risk Management of Digital Information: A File Format Investigation" (http://www.clir.org/pubs/reports/pub93/contents.html)

[8] See the complete workshop report at: http://www.si.umich.edu/digarch/Report.DFt.2.doc

[9] The question can also be posed from a different perspective: given the vast corpus of information in digital form, what information are we *not* going to preserve, or in other words, what information are we willing to put at risk? The underlying assumption here is that the rate at which digital information is created exceeds our ability to preserve it all, although this view is not universally held. But if preservation resources are scarce enough that choices must be made, as seems likely, then a decision to preserve a certain set of materials is synonymous with a decision not to preserve the rest. This leads into the issue of *selection*, which is a topic of great interest in the digital preservation community.

preservation is an economically sustainable process. Much of the existing literature associated with the economics of digital preservation tends toward this direction.

A reckoning of the costs and revenues associated with a particular digital preservation activity is, of course, a key component of the economics of digital preservation. It is a necessary ingredient for evaluative tools such as cost/benefit analysis, which serve to define precisely the *incentives* for participating in an economic activity. Cost/benefit analysis will, from time to time, reveal cases where sufficient incentives fail to appear: i.e., the costs outweigh the benefits. How does digital preservation fare in this regard? Unfortunately, measuring the costs and benefits associated with digital preservation can be exceedingly difficult, a point which will be discussed later in the paper. By extension, then, it is often difficult to predict *a priori* whether or not a particular digital preservation activity is economically viable.

A further complication is that digital preservation can take place under a variety of scenarios, the specific circumstances of which have a profound impact on the nature and extent of the incentives to preserve. For example, some organizations are interested in the preservation of proprietary materials; others are concerned with the preservation of materials in the public domain. Some digital preservation initiatives view economic issues in light of setting up and managing a repository internal to an organization; others adopt the perspective of a centralized repository providing digital preservation services to multiple organizations. Some digital repositories are created to accommodate materials whose deposit is required; others are aimed at providing digital preservation services to organizations who are self-motivated to preserve digital content under their control. In short, there is no universal characterization of the incentives to preserve; they can be discussed only superficially without reference to the scenario under which they emerge.

Given these considerations, there seems little promise that a systematic study of the incentives to preserve digital materials would yield anything of general applicability. But in fact, powerful insights into the nature of these incentives can be obtained through the use of a fairly terse framework that distills the digital preservation process down into a few key decision-making roles. These in turn serves as the basis for defining a small set of core scenarios under which digital preservation might take place. With this framework in hand, application of a little economic intuition serves to identify which of these scenarios are most likely to fail to yield sufficient incentives to carry out digital preservation. These insights can be obtained even without the aid of formal cost/benefit studies that precisely quantify the incentives to preserve.

Economic theory embodies the fundamental assumption that people do things – buy a house, eat at a restaurant, invest in the stock market – not simply because they can, but because they have a compelling interest in doing so. So it is with preservation, as the example of the Bermuda Hundred battlefield sites illustrates. In that case, a socially desirable end – preservation of part of the cultural record – was hindered by the fact that those who controlled the objects to be preserved had little or no incentive to participate in the preservation process. Cases where an economic activity is generally considered desirable, but will not be carried out through the voluntary transactions of privately

motivated entities, are examples of *market failure*. A great deal of economic theory is devoted to identifying situations where appropriate incentives fail to appear, why they fail to appear, and what remedies can be employed to correct for this failure. These theories can be usefully applied to the core scenarios of digital preservation suggested by the simple framework described above.

Understanding the characteristics of the scenarios where the incentives to preserve digital materials are weak or non-existent is an important first step in addressing the economic challenges of digital preservation. It is a necessary foundation for more focused work aimed at developing cost and/or business models for specific digital preservation initiatives. It also informs the development of remedies to correct for a lack of incentives in situations where this is appropriate. Finally, it contributes toward an understanding of the prospects for developing market-based digital preservation services.

* * *

In their book on the economics of the new Information Economy, Carl Shapiro and Hal Varian argue that the New Economy did not imply a new economics.[10] The same time-tested principles that guided the development of the "Old Economy" were still fundamental to understanding the new one. In the same way, no new economics is needed to understand the fundamental economic issues associated with digital preservation; existing theories of economic behavior, many of them quite old, can shed a great deal of light on the economic issues surrounding the relatively new topic of ensuring the long-term retention of digital materials.

This paper provides an overview of economic issues related to the incentives to preserve digital materials. Section II discusses challenges and limitations in enumerating and quantifying costs and revenues in the context of digital preservation. Section III discusses the more fundamental issue of the incentives to preserve, and describes two key facets to this concept. Section IV develops a simple framework for analyzing the incentives to preserve, based on three principal decision-making roles among stakeholders in the digital preservation process. These roles are combined and re-combined into a set of "core scenarios" under which digital preservation activities might occur. Section V offers economic explanations for the nature and extent of the incentives to preserve suggested by the characteristics of each scenario, as well as possible remedies when these incentives prove to be insufficient. Section VI enumerates a list of topics for future research in the economics of digital preservation, based on the foundations established in this paper.

II. COSTS AND REVENUES

It is useful to begin with a brief survey of existing work relevant to the economics of digital preservation. Much of this work concerns the enumeration and quantification of costs, as well as the development of cost-recovery or profit-making mechanisms to

[10] Shapiro, C., and H. Varian (1999) *Information Rules: A Strategic Guide to the Network Economy* (Harvard Business School Press, Boston)

support the provision of digital preservation services. While some progress has been made in these areas, cost and revenue models for digital preservation are notoriously difficult to quantify and resistant to generalization. As such, they fall short in terms of adequately characterizing the incentives to preserve digital materials.

* * *

The 1996 report of the Task Force on Archiving of Digital Information is widely referenced as the authoritative reckoning of the challenges posed by digital preservation. "In addition to managing their operating environment and the migration of information through hardware and software platforms," the Task Force observes in this report, "a third function by which digital archives fulfill their commitment to preserve electronic information is in managing the costs of these activities."[11] Other sources have echoed this message: economic issues, as well as technical ones, must be confronted in the course of securing the long-term integrity and accessibility of digital materials. The most elegant technical solutions are of little more than academic interest if the cost of implementation is prohibitive or unrecoverable.

The Task Force report identified two areas of focus relevant to the economics of digital preservation:

- **Cost Modeling:** "… the detailed interplay of cost factors that might promote the kinds of specialization, division of labor, and competition needed … to drive digital archives not just to manage costs against a standard of information quality and integrity, but to strive vigorously to lower those costs while maintaining and improving the standard of quality."[12]
- **Financing:** "… how to balance [the costs of digital archiving] with income, either from a sponsoring organization or philanthropy that absorbs the costs or from direct or indirect charges for use."[13]

Since publication of the Task Force report, a significant literature has coalesced around the topics of cost modeling and financing for digital preservation.

II.1. Costs

The National Library of Australia's Preserving Access to Digital Information (PADI) Web site[14] provides a useful bibliography of the literature on costs. Much of this work falls into two broad categories: models describing the key components of the long-term costs of preserving digital materials, and case studies providing data from particular projects and initiatives. In regard to the first category, the most ambitious effort is Hendley's JISC study[15] on the cost implications of various digital preservation strategies.

---

[11] Task Force on Archiving of Digital Information (1996): "Report of the Task Force on Archiving of Digital Information", p. 30. See the complete report at http://www.rlg.org/ArchTF/

[12] Ibid, p. 31

[13] Ibid, p. 37

[14] http://www.nla.gov.au/padi/topics/5.html

[15] Hendley, T. (1998) "Comparison of Methods and Costs of Digital Preservation". Available at: http://www.ukoln.ac.uk/services/elib/papers/tavistock/hendley/hendley.html

Hendley's cost model is based on the seven stages, or modules, associated with digital collection management proposed by Beagrie and Greenstein.[16] Hendley expands this framework by enumerating a set of cost elements for each module, in order to describe the total costs associated with the long-term management of a digital collection. With this in hand, he then identifies the subset of costs directly or indirectly related to preservation. Hendley applies this general cost model to four categories of digital resources. Other attempts to identify cost elements associated with digital preservation include Russell and Weinberger[17], who organize cost elements according to the chronology of a collection manager's workflow; and Sanett[18], who develops a cost model for the preservation of authentic electronic records.

But a number of factors conspire to complicate analysis of digital preservation's costs. First, the availability of cost data from ongoing digital preservation activities is scarce. Much of this problem can be ascribed to the fact that the notion of preserving digital materials is relatively new: consequently, most digital preservation initiatives are still in their infancy and are as yet unable to supply a quantitative record of the full cost trajectory associated with maintaining digital materials over long periods. It is not surprising, then, that the bulk of the actual cost data available is confined to the earliest stages of the digital preservation process, such as selection, reformatting or conversion, and storage technologies. The greater part of the costs, however, are distributed over an indeterminate time horizon, hidden in the processes that ensure the long-term recoverability of the object's bit stream, as well as the ability to render the bit stream into useable information.

In the absence of empirical data, future costs of digital preservation can be estimated, but the uncertainties of long-term preservation impart a high degree of variance to these estimates. One source of uncertainty arises from the very *raison d'etre* for digital preservation: the constantly evolving technological environment in which digital materials exist. This creates an ever-present risk that archived digital objects will be "orphaned" when the hardware and/or software needed to store, render, or utilize them becomes obsolete or unavailable. The time horizon within which technological obsolescence becomes a concern can be uncomfortably short, or even surprisingly long. It is difficult to anticipate, *a priori*, the future demise of current hardware and software environments, and therefore, difficult to forecast the timing, frequency, complexity, and ultimately, the cost of intervention to stave off obsolescence.

A second source of uncertainty associated with estimating the costs of digital preservation is that these costs are a function of many variables, including, but not limited to, the period of archival retention (ten years? one hundred years? "perpetuity"?); storage technologies, including hardware and digital media; the level of access; the objectives of

[16] Beagrie, N. and D. Greenstein (1998) "A Strategic Policy Framework for Creating and Preserving Digital Collections". Available at: http://ahds.ac.uk/strategic.pdf
[17] Russell, K. and E. Weinberger (2000) "Cost Elements of Digital preservation". Available at: http://www.leeds.ac.uk/cedars/documents/CIW01r.html
[18] Sanett, S. (2002) Toward Developing a Framework of Cost Elements for Preserving Authentic Electronic Records into Perpetuity" *College and Research Libraries* Vol. 63 No. 5 p. 388-404

preservation; preservation strategies (migration? emulation?); type and variety of digital formats; richness of metadata description; and the relative mix of labor-intensive and automated processes. Accurate cost estimates must be predicated on a reasonably specific description of the preservation processes involved. This in turn suggests that there is little scope to generalize these estimates over heterogeneous digital preservation activities. In the context of institutional repositories, a recent SPARC report noted that these repositories "have been diverse in scope and varied in their technical execution. Such disparate experiences make it difficult to formulate a universal economic model and project either development or operating budgets for new institutional repositories."[19]

Finally, cost implications for digital preservation are difficult to characterize because the process of digital preservation is itself not yet well-understood. The menu of preservation strategies from which digital repositories can choose is generally perceived to include three alternatives: technology preservation, format migration, and emulation. But much work remains to be done in exploring alternatives for the long-term retention of digital materials. The strategies for digital preservation considered feasible today fall short of what may be termed mature, well-defined processes, and for the most part they remain untested in production environments. There is no assurance, therefore, that the current state-of-the-art digital preservation strategies will not themselves become obsolete tomorrow. The fact that the art and science of digital preservation is still in flux, with the likely consequence of future changes in accepted methods for preserving digital materials, introduces yet another source of uncertainty into estimating preservation costs.

II.2. Financing

The second area of focus identified by the 1996 Task Force report is financing. The issue of financing digital collections has received some attention, although little of this work focuses on the preservation component of the digital lifecycle. For example, the Pricing Electronic Access to Knowledge (PEAK) project[20] conducted experiments analyzing alternative pricing plans for access to electronic journals. An article by Sairamesh, et al.,[21] develops a fairly comprehensive economic framework for commercial digital libraries, focusing on pricing strategies for services, and mechanisms for charging and billing users.

More recently, CLIR published the proceedings of a conference aimed at identifying sustainable business models for digital collections maintained by libraries and museums.[22] Of particular interest was the business model followed by JSTOR, a digital archive for the scholarly journal literature. In a nutshell, JSTOR "acquires rights from publishers to full runs of selected journals in the humanities and social sciences, digitizes the content, and makes it available on the Web through institutional site licenses."[23] The

---

[19] Crow, R. (2002) "The Case for Institutional Repositories: A SPARC Position Paper", p. 27. Available at: http://www.arl.org/sparc/IR/ir.html

[20] http://www.lib.umich.edu/retired/peak/

[21] Sairamesh, J., C. Nikolaou, D.F. Ferguson, and Y. Yemini (1996) "Economic Framework for Pricing and Charging in Digital Libraries" *D-Lib Magazine* Vol. 2, No. 2

[22] "Building and Sustaining Digital Collections: Models for Libraries and Museums" (2001), available at: http://www.clir.org/pubs/reports/pub100/pub100.pdf

[23] Ibid, p. 3

JSTOR strategy for financing digital preservation is to capitalize the entire future stream of digital preservation costs into a single, upfront fee. This Archive Capital Fee (ACF) applies to access to a single collection within JSTOR's offerings. For example, access to JSTOR's Arts & Sciences I Collection for a medium-sized US academic institution is $4,000 a year. In addition to this annual fee, the institution must pay a one-time ACF of $25,000.[24]

Discussion of financing issues tend to focus on digital collections as access mechanisms, rather than as archival repositories. In other words, financing is based on the idea that users are primarily paying for access services, rather than preservation services. Even JSTOR, a self-described "archive", notes "[w]hile JSTOR's ultimate mission is to change ways of thinking about the economics of archiving in a digital realm, it remains a challenge for JSTOR to build trust as an archiving agent in the library world. Most libraries see it as a great delivery system, not an archive …"[25] There is little discussion of financing or pricing strategies for *digital preservation services alone*; rather, cost recovery or revenue-producing schemes usually focus on other digital library services.

One exception is a collaboration between Yale University Library and Elsevier Science exploring the preservation of large collections of commercially published electronic scholarly journals. A recent report on the collaboration's findings enumerated a list of alternative funding models specifically aimed at digital archiving[26]:

- **Up-front Payment:** Users pay for a defined quantum of storage and with that one-time payment comes an eternity of preservation. The up-front payment would be invested partly in ongoing archival development and partly in an "endowment" or rainy day fund.
- **Ongoing Archival Fees:** An "insurance premium" that gives an ongoing supply of money, adjustable as costs change, and modest at all times. This reduces the risk for the provider but increases the risk for the beneficiary.
- **Traditional Library Model:** Library (or museum or archive) picks up the tab and is funded by third-party sources.
- **Fee for Services Operation:** The archive provides certain services (special metadata, support for specialized archives) in return for payments.
- **Hybrid:** If no single arrangement seems sufficient then a hybrid likely will emerge, perhaps with one set of stakeholders sharing the up-front costs, while another enters into agreement to provide ongoing funding for maintenance and potential access.

The "traditional library model" is frequently encountered in digital preservation activities currently taking place in the cultural heritage community, although the up-front model has also been utilized (e.g., in JSTOR). Most of the remaining models remain largely unexplored.

---

[24] See http://www.jstor.org/about/asI.fees.html
[25] "Building and Sustaining Digital Collections: Models for Libraries and Museums", p. 5
[26] "The Yale Electronic Archive: One Year of Progress" (2002), available at: http://www.library.yale.edu/~okerson/yea/ p. 30

* * *

In summary, the literature associated with the Task Force's two focus areas – cost models and financing – is significant, but still a work in progress. As the SPARC report notes, "[t]he cost of preservation archiving, for any digital collection, has yet to be definitively determined."[27] Some broad descriptions of the major cost elements of digital preservation have emerged, along with a small number of case studies providing real-world cost data. Some work dealing with financing or cost recovery issues has also been conducted, but, as noted above, the focus of this literature is not on financing associated with the provision of digital preservation services. It is fair to say that the type and magnitude of the long-term costs of digital preservation, as well as the best means to recover these costs, are still open questions.

III. INCENTIVES TO PRESERVE

Organizations faced with the prospect of preserving digital information are naturally keen to understand the costs, and possibly revenues, associated with this activity. Cost and revenue models are therefore essential for developing sustainable digital archives. However, they embody the implicit assumption that digital preservation is in fact going to take place. For example, a cost model, taken by itself, simply describes the type of costs that will be incurred *if* a certain digital preservation activity is undertaken. But the fundamental question is, of course, *will* digital preservation be undertaken? To answer this question, the cost and revenue models must be quantified as a prelude to some form of cost/benefit analysis. The difficulties with quantification were documented in the previous section. Consider, then, a body of digital materials deemed to be "at risk". There seems to be a scarcity of predictive tools that shed light on the question of whether any stakeholder will, or even should, undertake to preserve these materials.

In the absence of cost/benefit analysis based on precise quantitative data, it would be useful to have a qualitative model offering some predictive capacity in regard to the question *will digital preservation take place?* Fortunately, economic theory provides a number of tools useful for this purpose. In particular, a great deal of general insight into the question can be obtained from analyzing a single economic concept: the incentive to preserve.

The incentive to preserve has two facets:

- **The Incentive to Preserve (1):** perceived motivation sufficient to induce a party to recognize a need to take action to secure the long-term viability of digital materials in which they are a stakeholder.
- **The Incentive to Preserve (2):** perceived motivation sufficient to induce a party to develop and implement technologies aimed at ensuring the long-term viability of digital materials.

[27] Crow, p. 28

A loose analogy can be drawn between the two facets of the incentive to preserve on the one hand, and a market for digital preservation services on the other. The first incentive produces what might be termed the *demand side* of the market: in other words, the collection of entities (individuals, businesses, institutions) that identify a need to effect the long-term retention of digital materials under their ownership or control, or in which they are stakeholders, through the purchase of some form of digital preservation service. Two primary factors can motivate this perceived need. First, these entities may benefit directly from long-term preservation – as a user of the materials, for example, or through the fulfillment of evidentiary or legal obligations to retain the materials over time. Second, a need to preserve may arise in conjunction with, or as a component of, a service or function the entity provides on behalf of some separate constituency: e.g., a commercial service provided by a for-profit entity to fee-paying customers, or a public service provided by a mission-driven institution to society at large.

The second incentive produces what can be termed the *supply side* of the market: i.e., the collection of entities that are willing to meet the needs expressed by the demand side of the market, by implementing and making available services to effect the long-term preservation of digital materials. The primary factor motivating this willingness is the perception that the capacity to provide digital preservation services would contribute toward the fulfillment of some over-arching objective governing the entity's activities: e.g., to increase the value of the entity through the accumulation of revenues or profits, and/or to allow the entity to achieve its stated mission or function.

It should be noted that likening the incentives to preserve to the demand and supply sides of a market does not imply that the motivations underlying digital preservation are strictly pecuniary in nature. Incentives to preserve digital materials in the first sense (the "demand side") can arise from a variety of sources, ranging from a desire to safeguard the condition of a set of digital materials in anticipation of future economic value, to an altruistic desire to ensure the availability of digital materials for research and education. Similarly, incentives to preserve in the second sense (the "supply-side") can range from a perceived opportunity to generate revenues, to a perceived responsibility to perpetuate the cultural record.

The incentives to preserve are straightforward to define conceptually, but often difficult to apply in practice. Several factors serve to muddy the waters. First, while the framework of a market is helpful in distinguishing and describing the two incentives, digital preservation does not always occur in the form of an arms-length market transaction. Oftentimes, the entity demanding digital preservation services and the entity supplying them are one and the same. This is not to say that in such circumstances, decision-making proceeds independent of an evaluation of incentives. Rather, these incentives are not made explicit in the form of a two-party market transaction. This issue will be discussed further in the next section.

Second, the nature and magnitude of the incentives to preserve are not necessarily static, but rather, may evolve over time. For example, over the short- or even medium-

term, commercial publishers may perceive sufficient incentives to take action to preserve their digital content in order to protect existing revenue streams, or to safeguard opportunities to earn future revenue streams. During this time, the preservation needs of publishers could be similar to those expressed by other stakeholders in the long-term retention of the materials, such as libraries. But in regard to the very long-term[28], when the economic value of the materials has perhaps been exhausted, the incentives to preserve, from the publisher's perspective, may disappear; in this event, a scenario emerges where there is a gap between preservation objectives (e.g., the need to preserve the cultural and scholarly record in perpetuity) and preservation incentives on the part of relevant decision makers (in particular, the publishers, who are the owners of the digital materials). The temporal aspect of preservation incentives has been recognized in practice. Libraries and other subscribers/licensors of digital content fear that in the absence of measures taken to achieve long-term preservation, expensive investments in digital content could be jeopardized. But a joint statement by the International Federation of Library Associations and Institutions (IFLA) and the International Publishers Association (IPA) in 2002 observed that "while publishers generally can ensure the short-term archiving of their publications so long as these publications are economically viable, libraries are best-placed to take responsibility for long-term archiving through appropriate arrangements with publishers."[29] The relationship between time and the incentives to preserve will be re-visited later in the paper.

A third difficulty emerges from the fact that the decision to preserve materials, digital or otherwise, is sometimes guided by institutional mission, public policy, or legal requirement, rather than free economic choice. For example, the US National Archives and Records Administration (NARA) is legally obligated to preserve the records of government agencies that are transferred into NARA's custody; similarly, banks must retain records of financial transactions for a legally prescribed period. In these situations, a discussion of incentives is of little practical interest. This is not to dismiss or minimize the economic issues faced by "mission-driven" institutions where digital preservation is a vital component of their stated function, or businesses that must preserve digital materials to meet legal or evidentiary requirements. But in cases such as these, the fundamental economic issue is not one of incentives, but more likely one of *opportunity cost*.[30] At the "macro" level, this might take the form of deciding which institution or organization is best placed to fund and/or carry out a particular digital preservation activity; at the "micro" level, it may be a question of re-allocating budgeted funds to accommodate new digital preservation responsibilities.

But in scenarios where digital preservation is a question, rather than an assumption, the fundamental economic issue is whether or not relevant decision-makers will choose to allocate resources to this activity. It is in this context that incentives assume a role of critical importance. Many, if not most, of the economic issues debated within the digital preservation community today can ultimately be traced back to the

[28] Admittedly, a very vague time frame.
[29] http://www.ifla.org/V/press/ifla-ipa02.htm
[30] Opportunity cost is the benefit foregone by choosing to allocate a set of resources to a given activity, rather than to the next best alternative use.

motivations, or lack thereof, of stakeholders to participate in the preservation of digital materials, either through the identification of a *need* to preserve, a *willingness* to supply digital preservation services, or both. In these circumstances, the following economic question can be posed:

> *Do sufficient incentives exist for relevant decision-makers to, on the one hand, identify a need to take action to preserve a given set of digital materials, and on the other, provide digital preservation services to parties interested in utilizing them?*

It is thesis of this paper that this is the fundamental economic question associated with digital preservation.

IV. A SIMPLE ECONOMIC FRAMEWORK

Digital preservation takes place in a variety of circumstances. An organization can undertake to archive its digital materials internally, or it can contract with an external party to perform this service. The purpose of archiving can be to ensure the long-term retention of proprietary business records to meet evidentiary requirements, or it can be to preserve important cultural artifacts for future generations. Archiving can be compulsory, as in the case of meeting legal requirements or organizational obligations, or it can be a strictly voluntary activity. Given the multiplicity of scenarios under which digital archiving might occur, the incentives to preserve cannot be universal in character.

Unfortunately, in workshops and conferences where the economic issues of digital preservation are discussed, participants in the dialog often represent a wide cross-section of these archiving scenarios. While consensus can be reached at a superficial level (e.g., economic issues are important, organizations of all varieties must confront them, there is a critical need for research in this area), there is difficulty in moving to more detailed discussion, as the specific circumstances of various digital preservation scenarios create multiple points of departure for a discussion of this kind.

How then to correlate the variety of circumstances under which digital preservation takes place with the incentives to preserve digital materials? The methodology of economics helps to shape the problem. Economics is a behavioral science: that is, it describes and analyzes the behaviors of individuals or organizations of individuals engaged in economic activities. For example, economic theory identifies certain key entities whose activity, in aggregate, constitutes a functioning economy – firms, consumers, labor, government, etc. Economic analysis of digital preservation, therefore, should describe and analyze the behaviors of individuals or organizations of individuals engaged in the activity of preserving digital materials.

A necessary first step, then, in analyzing the incentives to preserve is to identify the key decision-makers whose behavior these incentives are expected to impact. In this section, a simple framework is developed that 1) identifies the key decision-making roles associated with digital preservation, and 2) organizes these roles into a set of "core

scenarios", each reflecting a particular organizational structure for digital preservation activities.

IV.1. Key Roles in the Preservation of Digital Materials

Economic analysis of a market examines the behaviors of two highly stylized economic roles: the consumer and the firm.[31] The consumer is the decision-maker on the demand-side of the market: he or she determines whether or not to purchase the good or service offered in the market. If the decision is in the affirmative, the consumer reaps whatever benefits are associated with ownership of the purchased good. The firm, on the other hand, is the decision-maker on the supply-side of the market: it determines whether or not to produce a particular good or service, and in the event production takes place, garners revenues from the sale of the good in the market.

It is straightforward to see how incentives work in this simplistic framework. Consumers participate in the market when they perceive that the benefits from purchasing the good exceed its cost. Firms participate in the market when the market price of the good exceed production costs. The co-existence of these incentives results in the development of a self-sustaining market from which all participants derive mutual benefit.

Digital preservation presents some interesting twists to this basic paradigm. Two are of particular importance. First, digital preservation is a service that is used in conjunction with another good – specifically, a digital resource. Purchase of the digital preservation service is intended to enhance the value of the digital resource, primarily by extending its useful life.[32] Second, the entity that benefits from digital preservation is not always the entity that owns the digital resource. These two factors profoundly impact the nature and extent of the incentives to preserve. This point will be discussed in detail throughout the remainder of this paper.

What, then, are the key economic decision-making roles in digital preservation? Three principal roles[33] are evident among stakeholders in the digital preservation process:

Rights Holder: holds the intellectual property rights to one or more digital resources. More specifically, the Rights Holder is the entity that possesses the legal authority to 1) preserve the materials, and 2) provide access to these materials during the period of their

[31] In some economic models, the decisions of other entities are modeled as well: for example, the labor and capital owners in the factor markets, or government regulators.

[32] The idea that preservation enhances the value of a digital resource must, of course, be predicated on the assumption that the resource is perceived as being valuable, or alternatively, that it is likely to be valuable in the future. The latter case may be referred to as "hypothesized value", a term for which the author credits and thanks Paul Conway. "Value" is here used in its broadest sense, encompassing economic value as well as cultural or scholarly value. In short, a digital resource must be perceived as being valuable, either now or in the future, for any investment in digital preservation to be justified economically. As will be seen below, this is akin to assuming the presence of a Beneficiary associated with the digital preservation process.

[33] The author thanks Donald Waters for his helpful comments on the nature of the decision-making roles associated with digital preservation. For a different, yet related perspective on roles, organizational models, and their implications, please see Mr. Water's forthcoming paper "Doing Nothing is Not an Option: Evolving Roles and Opportunities in the Digital Era".

retention in an archive, and 3) grant or cede the rights to preserve or provide access to another entity.

Defining the Rights Holder as a single, monolithic role is problematic, because in practice, intellectual property rights take the form of a "bundle" of rights, which can be split apart and distributed among multiple entities. For example, under the principal of legal deposit, a national library may have the right to preserve a copy of a published item, while the publisher retains other rights, such as the right to reproduce the materials, or to exploit them commercially. For the purposes of this analysis, the Rights Holder is understood to mean the entity upon whom the rights bundle is initially conferred (e.g., at the time of the digital object's creation), and who subsequently has the authority to dispose of these rights – either by retaining them or granting them to another entity – as they see fit.

Beneficiary: derives some form of benefit from the long-term retention of digital materials. The Beneficiary role includes entities that benefit directly from preservation – i.e., the end users or consumers of the preserved materials – as well as those who serve an intermediary function (such as libraries and other cultural heritage institutions) in terms of facilitating end users' access to digital materials. In considering incentives to preserve, however, the relevant entity is the one in which the decision-making ability is vested – in particular, the ability to identify or select which digital materials warrant preservation. In some scenarios, the end users may make this determination, but often it is the intermediary institution, acting on behalf of (and perhaps in consultation with) the direct beneficiaries, to which this decision is entrusted.

The first facet of the incentive to preserve is associated with the Beneficiary role.

Archive: implements and manages the processes and systems needed to achieve the long-term preservation and accessibility of digital materials. The role of the Archive is to secure the long-term retention of digital materials placed in its custody. In this sense, the focus of this role is on the preservation process itself. It does not include the process of identifying a need to preserve a certain set of materials – in other words, the process of selection. This function is subsumed within the Beneficiary role.

The second facet of the incentive to preserve is associated with the Archive role.

There is a correspondence between the roles defined above, and those identified in the Open Archival Information System (OAIS) reference model[34] in its description of the external environment in which an archive operates. Here, the notion of "external" is applied in the context of the archiving system itself, rather than in terms of an organization or institution. In addition to the archiving system, the OAIS reference model identifies three other roles in the external environment:

[34] The reference model is available at: http://wwwclassic.ccsds.org/documents/pdf/CCSDS-650.0-B-1.pdf

- **Producer:** the role played by those persons, or client systems, who provide the information to be preserved[35]
- **Consumer:** the role played by those persons, or client systems, who interact with [archive] services to find preserved information of interest and to access that information in detail[36]
- **Management:** the role played by those who set overall [archive] policy as one component in a broader policy domain[37]

Producers are those who deposit information into the archiving system to be preserved; in the context of the economic roles, they correspond to the Rights Holders. Consumers are those who access and utilize the information preserved in the archiving system; they correspond to the economic role of Beneficiaries. Finally, the OAIS roles of Management (responsible for guiding the archive's overall direction and policies) and the archiving system itself are combined into one economic role: Archive. Despite this correspondence, there are subtle yet important differences between the OAIS roles, and the economic roles defined above. The OAIS roles are oriented toward a *functional*, or workflow view of an archiving environment. The economic roles, on the other hand, are defined to emphasize the key decision-making components of the archiving environment.

Note that that in the context of both the OAIS and the incentives to preserve, we are considering roles, rather than distinct entities: a single entity can subsume more than one (or even all) of the roles described above. Note also that reference to any role in the singular (e.g., Rights Holder, rather than Rights Holders) does not necessarily imply that only a single entity can assume that role. In many cases, multiple entities can simultaneously serve in any one of these roles.

The three economic roles of Rights Holder, Archive, and Beneficiary are sufficient to analyze the incentives associated with digital preservation. As will be shown later in the paper, the source of many of digital preservation's economic challenges can be found in the manner in which these roles are distributed among one or more distinct entities, and the resulting impact on the underlying incentives to preserve.

IV.2. Organizational Models for Digital Preservation

On the face of it, the mapping between decision-making roles and incentives seems straightforward: the incentive to demand digital preservation services should emerge from the Beneficiary (who benefits from the services), and the incentive to supply digital preservation services should emerge from the Archive (whose existence is predicated on providing these services). But what of the Rights Holder, who controls the disposition of the digital materials? As will be made clear below, the identity of the Rights Holder – in particular, whether the entity assuming this role also assumes the role of Archive, Beneficiary, both, or neither, is a critical factor in determining the existence of sufficient incentives to carry out digital preservation.

[35] OAIS reference model, p. 1-12
[36] Ibid, p. 1-9
[37] Ibid, p. 1-11

To see this, we begin by enumerating the various combinations of the three economic roles that can be assigned to one, two, or three distinct entities. Five different combinations are possible:

CENTRIPETAL MODEL …

… Rights Holder, Archive, and Beneficiary are the same entity

In this model, all decision-making roles are compressed within a single entity. The key features of this model are 1) that both facets of the incentive to preserve – the need to preserve a set of materials, and the willingness to carry out their preservation – are manifested in a single entity, and 2) the authority, or right, to preserve the materials in question is vested in this same entity.

CENTRIFUGAL MODEL …

… Rights Holder, Archive, and Beneficiary are all separate entities

In this model, all decision-making roles are spread out among separate entities. More specifically, one entity has identified a need to preserve a set of materials; another entity has expressed a willingness to provide services to effect the preservation of these materials; and a third entity holds the decision-making authority to initiate preservation.

SUPPLY-SIDE MODEL …

… Rights Holder and Archive are the same entity; Beneficiary is separate

In this model, the willingness to provide preservation services, as well as the authority or legal right to initiate preservation, are combined in a single entity. A separate entity embodies the need to preserve a set of digital materials. The key feature of this model is that one facet of the incentive to preserve – the willingness to provide (supply) digital preservation services for a set of digital materials – is reinforced by the legal authority to initiate preservation on the digital materials in question.

DEMAND-SIDE MODEL …

… Rights Holder and Beneficiary are the same entity; Archive is separate

In this model, the authority to preserve a set of digital materials, as well as the need to preserve these materials, is combined in a single entity. A separate entity embodies the willingness to provide digital preservation services. The key feature of this model is that one facet of the incentive to preserve – the need (demand) to preserve a set of digital materials – is reinforced by the legal authority to initiate preservation on the digital materials in question.

CONSOLIDATED MODEL …

… Archive and Beneficiary are the same entity; Rights Holder is separate

In this model, a single entity embodies both facets of the incentive to preserve: i.e., a need to preserve a set of digital materials, combined with a willingness to undertake the preservation of these materials through the provision of digital preservation services. The authority to initiate preservation, however, is vested with a separate entity. The key feature of this model is that although both facets of the incentive to preserve – the need to preserve a set of digital materials, *and* the willingness to preserve the materials – is manifested in the same entity, neither facet is reinforced by the legal authority or right to

initiate preservation. More specifically, the incentives to preserve, and indeed, the preservation process itself, are "consolidated" in a single entity, but the right or legal authority to initiate preservation is held by another entity.

The five models described above represent a fairly diverse range of organizational structures. At the endpoints are the extreme cases, in which the three decision-making roles of Rights Holder, Archive, and Beneficiary are either vested in a single entity (the Centripetal Model), or in three distinct entities (the Centrifugal Model). Interspersed between the endpoints are three intermediate cases (the Supply-Side Model, the Demand-Side Model, and the Consolidated Model), where the three decision-making roles are distributed in various combinations between two distinct entities. Each of these organizational models represents a unique combination of the three decision-making roles associated with digital preservation, producing a correspondingly unique interplay between the need to preserve a set of digital materials, the willingness to provide digital preservation services, and the authority, or legal right, to initiate the digital preservation process.

The organizational models should be regarded as describing the structure of a *particular* digital preservation activity, rather than the complete range of digital preservation activities in which a given entity participates. Libraries, publishers, archives, and other institutions can and do participate in multiple organizational models simultaneously.

One thing that is immediately apparent from the organizational models is the interplay between incentives on the one hand, and the legal environment in the form of intellectual property rights on the other. The prospects for digital preservation taking place are clearly improved when one or both facets of the incentive to preserve are amplified by the legal right to initiate preservation. This paper focuses on the economic aspects of digital preservation, which are embodied in the incentives to preserve. However, an important theme in the discussion that follows is that the incentives to preserve in a given scenario are heavily impacted by the manner in which the Rights Holder role is manifested.

The organizational models described above are simplistic, and real-world examples do not necessarily fall neatly into a single model. But the models provide sufficient context for considering issues associated with the incentives to preserve digital materials. For some of these models, the incentives appear to be well-defined: e.g., for the Demand-Side Model, it is easy to imagine an institution that owns digital materials having the incentive to take action to ensure their long-term viability, and on the other side of the coin, another organization willing to offer services to meet this preservation need. But in other models, incentives are less apparent: in the Supply-Side Model, the entity that benefits from the preservation of digital materials is different from the one that owns the materials, and it is the latter that embodies the role of providing digital preservation services. Do sufficient incentives exist in this case to ensure that digital preservation will in fact be carried out?

Economics has a rich literature treating the nature and extent of incentives in scenarios similar to those described by the five organizational models. Application of economic theory to the organizational models provides insight into 1) the prospects in each model for sufficient preservation incentives to emerge; 2) characteristics of each model that may diminish or eliminate these incentives; and 3) possible remedies for scenarios where sufficient incentives to preserve fail to emerge naturally. The next section is devoted to these questions.

V. AN ECONOMIC TOUR OF DIGITAL PRESERVATION

In thinking about digital preservation as an economic activity, or more specifically, the incentives for decision-makers to carry out this activity, it is useful to frame the issues in the context of a market for digital preservation services. A market is a mechanism for buyers and sellers to engage in *voluntary* transactions involving the supply and demand of a particular good or service. Key to the concept of a market is the idea that buyers and sellers participate not because they are forced to, but because they have a compelling interest – an incentive – to do so.

Do such incentives exist in the case of digital preservation? On the surface, the answer would appear to be "yes". Information in digital form is notoriously fragile, and without active maintenance, tends to quickly deteriorate or become obsolete. Digital preservation services offer a means to extend the productive life of valuable, often expensive, digital resources. This suggests that digital preservation embodies real benefits, and that owners of digital resources would be willing to purchase these benefits, in the form of digital preservation services, from reliable suppliers. Indeed, given that the risks associated with the long-term retention of digital materials are widely acknowledged, one might predict that a market for digital preservation services would be quick to develop and expand.

Curiously, this has not been the case. Although few would dispute that, for well-understood reasons, digital materials are at risk, a market to mitigate that risk is now only tentatively emerging. Nearly seven years after the release of the seminal Task Force on Digital Archiving report in 1996, digital preservation remains an activity that is, for the most part, conducted on an *ad hoc* basis, usually in the form of special initiatives or "demonstrator" projects.

Part of the explanation for this can be attributed to the fact that digital preservation is still an immature process. This is true from an economic as well as a technical standpoint. Considerable uncertainty surrounds the prospects for carrying out digital preservation in an economically sustainable way. From the perspective of some decision-makers, whether Rights Holder, Archive, or Beneficiary, the ill-defined nature of digital preservation as an economic activity presents considerable risk. As a consequence, they may feel that perceived benefits are not yet sufficient to justify the initial investment.

But even apart from this consideration, decision-makers may still find little incentive to commit resources to digital preservation. As with the case of the Civil War battlefields discussed at the beginning of this paper, a misalignment of incentives with the objectives of digital preservation can prevent the desired preservation activity from taking place. Understanding why incentives to preserve digital materials might fail to emerge, and how this can be corrected, is a critical economic challenge for the digital preservation community.

In the discussion thus far, digital preservation has been repeatedly characterized as an economic activity. But why this is so, and what this means precisely, has been treated only superficially. Examining the role of incentives in economic decision-making is facilitated by adding more context to the notion of digital preservation as an economic activity. Accordingly, the following sections introduce the concept of an *aftermarket*, and explain, by correlating the characteristics of digital materials to the principles of aftermarkets and the rationale for their formation, why digital preservation may be usefully characterized as an *aftermarket service*.

With the concept of aftermarkets in hand, the issue of incentives to preserve is re-visited, but this time in the context of the motivations to participate in an aftermarket for digital preservation services. The analysis is conducted by first identifying several key aspects of preserving digital materials, and then tracing their impact on the incentives to preserve in the context of the five organizational models described in Section IV. As the discussion will show, each of the organizational models harbor characteristics that could potentially diminish the incentives to preserve. This leads to *market failure*, in the sense that diminished incentives on the part of one or more of the entities embodying the three decision-making roles inhibits the provision of digital preservation services through voluntary market transactions. In scenarios where sufficient incentives to preserve are lacking, strategies for remedying the resulting market failure are discussed.

V.1. A Short Digression: Digital Information and Durability

Economists divide goods into two broad classes: *durable goods* and *consumption goods*. A durable good, once purchased, is expected to confer a stream of benefits to its owner over an extended period of time. This may be contrasted to a consumption good, in which the benefits from purchase are obtained as a one-time experience. For example, an automobile is a durable good, in that it is expected to be usable for many years. A bottle of wine, on the other hand, is a consumption good: once the wine is drunk, it is no longer available for future use. The automobile owner enjoys the use of the same vehicle for a lengthy period of time after purchase; to enjoy a continuous stream of benefits from wine, one must purchase a continuous stream of wine. Of course, reality blurs the delineation between durable and consumption goods: the automobile may be a "lemon" and break down the moment it is purchased, and leftover wine may be available for some time after the bottle is opened. In general, however, the distinction between durable and consumption goods is both valid and useful.

Is information a durable good or a consumption good? Some types of information seem to bear the characteristics of consumption goods: for example, the weather forecast

on the nightly local news is useful for a very limited duration. By and large, however, information is a durable good: once obtained, it can be utilized over and over again far into the future. For example, a book can be re-read, a movie re-watched, a music CD re-heard. For the purposes of the discussion in this paper, information in general, and digital information in particular, will be treated as a durable good.

In a sense, information is infinitely durable: for example, the information that the battle of Gettysburg was fought on July 1, 2, and 3 in 1863 will survive in human memory indefinitely. But if instead one considers a book on the battle of Gettysburg, the time horizon of the information's durability is shortened considerably: paper does not last forever. So while the book can be read and re-read over an extended period of time, eventually the process of obtaining benefits from the book (accessing the information) will become more difficult or even impossible: pages may become detached from the binding and lost, or the paper may become brittle with age and disintegrate, the infamous "yellow snow" known to librarians.

If one considers information not just as abstract intellectual content, but as a combination of the intellectual content and the physical "package", or medium, in which it is delivered and accessed, it is much more meaningful to talk about the durability of information. In particular, while an information package is expected to confer benefits for an extended period of time, it is not infinitely-lived. Put another way, the information package tends to depreciate over time, just like an automobile or a refrigerator. Depreciation takes the form of increasing difficulty, or even the outright impossibility, of obtaining benefits from use of the durable good.

V.2. Aftermarkets for Durable Goods

To combat the problem of depreciation in durable goods, and to maximize the length of time that the good remains useful, it is often necessary to perform some form of maintenance on the good. For example, an automobile owner must change the oil, inspect the brakes, have the engine tuned, replace the tires, etc. Alternatively, the owner of a durable good can often take actions which enhance the good's ability to confer benefits over time, over and above what was expected or agreed upon at the time of purchase. Continuing the automobile analogy, a car owner can purchase products or services, such as a leather steering wheel cover or fuzzy dice, which enhance the utility he or she derives from ownership of the car.

The market in which a durable good is initially purchased is called the *primary market*; the market for goods and services which maintain and/or enhance the ability of a durable good to provide a stream of benefits over an extended period of time is known as the *aftermarket*. Aftermarkets represent important segments of many industries – automobiles, appliances, and home construction are prominent examples. To reiterate, the common feature of these industries is that the goods they produce are expected to last for an extended period of time, but are subject to depreciation or some form of degradation during that period. There may also be opportunities to enhance the value of the good after purchase.

V.3. Digital Preservation as an Aftermarket Service

As with other durable goods, information has spawned an aftermarket of products and services to increase the length of its useful life and enhance its value. Recall that aftermarkets offer goods or services that serve to *enhance* or *extend* the stream of benefits conferred by a durable good. Given this, it would seem that a great deal of aftermarket services are associated with information. For instance, services are available to catalog or describe information resources in order to facilitate search and retrieval. Making information resources readily available through precise searching enhances the value associated with their usage.

Preservation services also constitute an aftermarket for information. Preservation micro-filming services, for example, extend the longevity of paper-based information resources by providing surrogates for long-term storage. Re-binding books with broken spines or loose pages also constitutes an aftermarket service, in that it extends the useful life of the good purchased in the primary market – the book – thereby extending the period for which the good will continue to confer benefits upon its owners.

Given that the fragility of information in digital form is widely acknowledged – in the words of the 1996 Task Force report, the life of digital information is "nasty, brutish, and short"[38] – it would seem that an aftermarket providing digital preservation services would have sprung up quickly to service the growing corpus of information produced in digital form. Without explicit intervention to maintain its viability, digital information can quickly deteriorate or suffer from technological obsolescence. Thus, the conveniences endemic to information in digital form – ease of access, precise duplication, immediate online delivery, etc. – are tempered by the corresponding need to take active measures to secure its long-term viability. An aftermarket in digital preservation services would meet this need.

V.4. Market Failure

This has not yet come to pass. Digital preservation services, such as OCLC's Digital Archive,[39] have only recently become available. For many organizations, digital preservation is done on an *ad hoc* basis, or not done at all. In general, digital preservation has yet to become a routine component of managing a digital asset's life cycle. The fact that digital preservation has not yet emerged as a common "aftermarket service" for digital resources seems incongruous with the status of digital information as a durable good. Aftermarket services which extend the life of other durable goods are commonplace – auto repair for cars, plumbers and carpenters for homes, etc. Should we not expect that similar markets will emerge for digital materials? What characteristics of digital preservation might impede the formation of sustainable aftermarkets in this activity?

V.5. Characteristics of Digital Preservation and the Impact on Incentives

Digital preservation embodies at least three characteristics which could potentially diminish the incentives for decision-makers to take the steps necessary to

[38] Task Force on Archiving of Digital Information (1996), p. 2
[39] http://www.oclc.org/digitalpreservation/services/archiving/digital/

secure the long-term retention of digital materials. This in turn would tend to suppress the development and expansion of an aftermarket for digital preservation services.

It should be noted that the three characteristics discussed below do not necessarily constitute an exhaustive list of all the factors that could potentially impact the incentives to preserve. There are other nuances associated with digital preservation that could act to impede its development as an economic activity. However, the following characteristics are particularly significant in terms of degree of impact, as well as relevance to a wide range of digital preservation contexts.

If sufficient incentives to preserve fail to emerge naturally, it is sometimes necessary to institute policies to create them artificially. Fortunately, economic theory has a great deal to say about how this might be done. The discussion to follow touches briefly on some of the ways policies can be devised to remedy common market failures associated with digital preservation.

V.5.a. Positive Externalities in the Digital Preservation Process

It was noted earlier that digital preservation embodies three important decision-making roles: Rights Holder, Archive, and Beneficiary. In practice, these roles can be distributed among one or more separate entities. Situations where the roles of Beneficiary and Rights Holder are manifested through two distinct entities are often associated with insufficient incentives to carry out preservation. Recall that this was so in the case of the Civil War battlefield sites described at the beginning of the paper: in this example, the Beneficiaries (Civil War historians and enthusiasts) were distinct from the Rights Holders (property owners), with predictable effects on the level of preservation activity eventually carried out.

This same factor often plagues efforts to preserve digital materials. It is not difficult to identify circumstances where those who would benefit from the preservation of a set of digital materials are not the same as those who own the rights to these same materials. If the Rights Holder possesses no compelling interest in ensuring the long-term viability of the materials in question, it is unlikely that they would be willing to bear the cost involved in conferring the benefits on others. Consequently, the sole decision-maker in a position to initiate preservation perceives little incentive to do so, and it is therefore probable that preservation of the digital materials will not take place.

The economic concept behind this misalignment of preservation objectives and incentives is that of a *positive externality*. An externality is a side effect or consequence of an economic activity that impacts a party external to that activity. Most people are probably familiar with the concept of a negative externality, which is an externality that takes the form of a harmful impact on the external party. The textbook example is a factory that, in the course of producing widgets, also produces smoke that soils the laundry hanging outside a nearby residence. In this case, an economic activity (producing widgets) imposes a cost on a party external to the activity (the damage inflicted on the laundry by the smoke) that is not taken into account by the party responsible for carrying out the activity (the factory). The effect of this unrecognized external cost is a higher

level of production activity (and thus more smoke) than what, from the viewpoint of society as a whole, would be considered optimal. Because the factory is not forced to internalize the cost associated with the damaged laundry, it has no economic incentive to reduce the level of the economic activity that is producing the harmful side effect.

A *positive* externality yields an effect that is precisely opposite. In this case, the economic activity produces a side effect that represents a benefit to a party external to the activity. For example, a logging company might plant trees to replace those it cuts for timber. This has the obvious benefit of replenishing the logging company's future timber inventories, but also has the benefit of improving air quality in the region. In deciding how many trees to plant, however, the logging company is unlikely to take this latter effect into account; it will therefore plant just enough trees to satisfy its anticipated future inventory needs, without regard to the fact that the more trees planted, the better the air quality. The consequence of this ignored supplemental benefit is a level of the economic activity (planting trees) that falls short of what society would view as optimal. The logging company lacks the appropriate incentives to increase its level of tree planting, because it does not explicitly capture the benefits from doing so, and is therefore unwilling to bear the associated costs.

In the same way, when the Rights Holder of a set of digital materials is an entity distinct from those who would benefit from the preservation of these materials, a positive externality emerges. The Rights Holder does not benefit from the preservation of the materials, so in evaluating the economic pros and cons of this activity, they are unlikely to take into account the benefits preservation would confer upon others. As with the case of the logging company planting trees, the result is a less-than-optimal level of digital preservation activity – or in the extreme, no preservation activity at all.

Returning to the Civil War battlefield example, preservation was not carried out because the Rights Holder – the owner of the property on which the entrenchments and fortifications lay – was not the Beneficiary of the preservation; rather, this was a separate (and rather limited) group of Civil War historians and enthusiasts who happened to be interested in this relatively obscure campaign. In this sense, there was no incentive for the Rights Holder to voluntarily bear the cost of participating in the preservation process, despite the benefits that would be realized by others. As a consequence, the presence of the positive externality often resulted in no preservation being carried out at all.

Examples of positive externalities are also evident in digital preservation. One example is particularly instructive in illustrating the effects of positive externalities on the incentives to preserve. Rather than purchasing digital content outright, libraries often enter into licensing or subscription agreements, by which the libraries' patrons gain access to the content, but ownership, and indeed, physical possession of the "bits", remains with the publisher. As owner of the content, the publisher fills the role of Rights Holder, with the legal authority to preserve the materials. It is likely that the publisher will make efforts to preserve the digital materials in order to maintain their accessibility as long as it is economically beneficial to do so. Once that threshold is passed, however, there is no guarantee that the publisher will continue to preserve the content, despite the

fact that libraries and their patrons would realize benefits from doing so. In making the decision not to take steps to preserve the content over the long-term – i.e., beyond what might be considered the digital materials' economic life – the publisher fails to take into account the full societal benefits from preservation, and therefore chooses a level of preservation activity less than what society as a whole might consider optimal.

In general, the presence of positive externalities, brought about when Beneficiaries are distinct from Rights Holders, tends to reduce or eliminate digital preservation activity when the externality is not explicitly accounted for in the incentives perceived by the Rights Holder.

How do positive externalities impact the development of an aftermarket for digital preservation services? Aftermarkets provide services which allow the owners of durable goods to extend the useful life of their property. The implication, of course, is that the owners of the property benefit from the purchase of aftermarket services. But, as the case of licensed or subscription-based digital content illustrates, the owners of digital materials are not always those who benefit from preservation. In these circumstances, the Rights Holder would not voluntarily participate in the aftermarket for digital preservation services. Beneficiaries, of course, would participate, but do not embody the decision-making role relevant to ensuring that these services are purchased and utilized.

Contrast the case of aftermarket preservation services for digital materials to those for paper-based materials. In the case of the latter, a market for preservation services has been established for some time, including such services as bookbinding and preservation micro-filming. But libraries and other institutions typically own paper-based materials outright, and are therefore free to contract for aftermarkets services to extend their longevity. In these circumstances, the benefits from aftermarket services accrue to the owners of the good, and, therefore, appropriate incentives exist for the owner to participate in the aftermarket.

In the preceding discussion, the underlying assumption is that the roles of Rights Holder and Beneficiary are held by distinct entities, with the implication that the Rights Holder perceives little or no benefit in preservation. But as mentioned earlier in the paper, the incentives to preserve can evolve over time. It may be that this perceived lack of incentives, and hence, the positive externality, could exist only in the long-term. For example, the Rights Holder may initially associate preservation with the ability to extract economic value (e.g., through licensing or subscription revenues) from a set of digital materials over the short- and medium-term. In this case, the Rights Holder would indeed perceive an incentive to preserve, and therefore embody a Beneficiary role. But once this economic value has dissipated, so too might the incentive to preserve. In this event, the Rights Holder may not find it worthwhile to continue preservation efforts. But this decision would be detrimental to other Beneficiaries: e.g., the preservation needs of libraries and other institutions acting as stewards of the scholarly and cultural record usually extend beyond the economic usefulness of the materials. The positive externality, and its impact on incentives, must once again be confronted.

Returning to the organizational models discussed in Section IV, it is straightforward to identify those which are most likely to embody insufficient incentives to preserve because of the presence of a positive externality. The positive externality occurs because the roles of Rights Holder and Beneficiary reside with two separate entities. This condition holds in three of the organizational models: the Centrifugal Model, the Supply-Side Model, and the Consolidated Model. In each case, the incentive to preserve is diminished because the parties who would benefit from digital preservation are entities distinct from those who control the long-term disposition of the digital materials in question.

In economics, the textbook solution for correcting a positive externality is to subsidize the activity in question. The subsidy artificially creates the benefit from performing the activity that the Rights Holder does not naturally perceive. Suppose it would cost the Rights Holder $100 to archive one gigabyte of data "in perpetuity". Then, a subsidy of $100 per gigabyte of archived data would at least make the Rights Holder indifferent between archiving or not archiving. Put another way, the subsidy is of sufficient magnitude such that the Rights Holder is made no worse off by performing the digital archiving activity.

Whether the subsidy would have to cover the full cost of preservation in order to create sufficient incentives to preserve for the Rights Holder depends on the particular circumstances. It may be that the Rights Holder, in an effort to cultivate goodwill with Beneficiaries, would be willing to assume part of the cost of digital preservation; in this case, only a fraction of the preservation cost would require subsidization. In general, the subsidy, in combination with any other non-economic incentives the Rights Holder perceives, considered in net of the costs associated with the digital preservation activity in question, should create sufficient incentive for the Rights Holder to initiate the activity. In doing so, the subsidy "internalizes" the spillover benefits associated with preservation into the decision-making process of the Rights Holder.

Who would be responsible for funding the subsidy? Subsidies are often provided by third-party entities, such as public agencies or philanthropic foundations, who intervene to fund activities when traditional market mechanisms fail. However, it is also possible that the Beneficiaries themselves will subsidize the Rights Holder, perhaps through the payment of an archiving fee as part of a broad service agreement covering both access and long-term retention. In either case, the subsidy serves to compensate the Rights Holder for undertaking an activity for which it perceives no private incentive.

A subsidy paid to the Rights Holder to undertake preservation fits the scenario where the roles of Rights Holder and Archive are combined into a single entity, as in the case of the Supply-Side Model. But in the Centrifugal and Consolidated Models, the Archive role is filled by an entity distinct from the Rights Holder. The subsidy would then be earmarked for the entity that fills the Archive role. In these circumstances, the primary issue is not direct subsidization of the Rights Holder to fulfill the Archive role, but rather to induce the Rights Holder to cede or grant the right to preserve to another entity who will perform this function.

This could be accomplished through a bargaining process between the Rights Holder, who holds the right to initiate preservation, and the Beneficiary, who benefits from preservation. In this scenario, there is a *cooperative surplus* from preservation – a net benefit from the activity that would make at least one of the parties better off, without making the other party worse off. For example, suppose again it would cost $100 per gigabyte to preserve digital content in perpetuity. Suppose further that if a gigabyte of data were preserved, the gain to Beneficiaries could be quantified at $175. In this case, there is a cooperative surplus of $75 associated with the activity: once the preservation activity has been fully funded (at a cost of $100), $75 in net benefits would remain.

The allocation of these net benefits is determined by the relative bargaining power of each party to the activity. If the bargaining position of the Rights Holder is particularly strong, for example, it may be that most of the benefits would accrue to that party: i.e., in order to receive the benefits from preservation of $175, the Beneficiary (or some third-party funder) would have to compensate the Archive $100 for carrying out the preservation, and then compensate the Rights Holder for ceding the right to preserve by allocating some portion of the net benefits to this entity – say $50. In this event, both parties are better off, in net, by undertaking the activity, but the Rights Holder captures the lion's share of the cooperative surplus.

Economic theory has little to say about who should garner the benefits in cases such as this. Suppose preservation is viewed as a legal right in the following context – from the Rights Holder's perspective, it is a right *not* to take any action to preserve digital materials under its control, while from the Beneficiary's perspective, it is a right to *require* preservation to take place at the Rights Holder's expense. If the two parties brought their case before a court to determine whose "right" should be upheld, how should the outcome be decided? A famous economic principle, known as the Coase Theorem, states that in the absence of transaction costs and other frictions in the bargaining process, it is irrelevant which party is granted legal entitlement: the efficient economic outcome will come to pass regardless. The logic is that the losing party will bargain with the winning party to achieve its desired outcome, re-allocating as much of the cooperative surplus as necessary to induce the latter to make the appropriate decision. For example, if it is makes economic sense for preservation to take place (i.e., benefits exceed costs), but legal entitlement is granted to the Rights Holder, the Beneficiary can bargain with the Rights Holder by allocating some portion of the net benefits to the latter as an inducement to preserve. Conversely, if preservation does not make economic sense (costs exceed benefits), but entitlement is granted to the Beneficiary, the Rights Holder can minimize its losses by paying compensation to the Beneficiary, up to the amount they would stand to lose if preservation was enforced, to induce them to give up their preservation objectives.[40]

[40] There are two other possible cases: 1) Preservation makes economic sense, and entitlement is granted to the Beneficiary: in these circumstances, no bargaining will take place and the Rights Holder will be forced to cover the expense of preservation, because by virtue of its legal right to have preservation take place, the Beneficiary has no need to share any portion of the net benefits with the Rights Holder in order to induce them to take action. 2) Preservation does not make economic sense, and entitlement is granted to the Rights

V.5.b. Digital Preservation as a Public Good

A second characteristic of digital materials is that in circumstances where multiple entities own the same digital resource (i.e., there are multiple Rights Holders), each of whom is also a Beneficiary in its preservation, the preservation of this digital resource takes on qualities of a *public good*. Public goods are a common source of market failure.

A discussion of public goods in regard to digital preservation must begin with a brief digression on uniqueness, or more to the point, *non-uniqueness*. Designation of a particular digital resource as unique or non-unique can be based on a number of factors, such as the intellectual content alone, or some combination of intellectual content and the "packaging" within which it is delivered (text formatting, file format, functionality, etc.) For the purposes of the discussion in this paper, a digital resource is considered *non-unique* when distinct copies are owned by multiple institutions, and these institutions perceive these copies to be acceptable substitutes.

The question of how non-unique digital resources impacts the aftermarket for digital preservation services proceeds from the economic concept of a public good. Donald Waters[41] offers an excellent discussion of how the cultural memory, and by extension, *preservation* of the cultural memory, are public goods. The analysis in this paper is an elaboration on his insights.

Public goods embody two important properties: non-rivalry and non-excludability. Non-rivalry implies that consumption of the good by one person does not diminish the amount of the good available for consumption by others. Non-excludability arises when "the cost of excluding non-paying beneficiaries who consume the good are so high that no private profit-maximizing firm is willing to supply the good."[42] A textbook example of a public good is national defense. The addition of one more citizen does not reduce the level of national security afforded to other citizens; also, if a citizen refuses to pay for national defense (say by refusing to pay his or her taxes), it is virtually impossible to exclude that one citizen from enjoying the protection from foreign invasion provided through the taxes paid by others.[43]

The combination of non-rivalry and non-excludability can lead to the failure of a market, composed of privately motivated individuals, to provide a socially optimal level, or quantity, of the good in question. This is caused by a lack of appropriate incentives. From the consumer's point of view, there is little incentive to purchase a non-rivalrous, non-excludable good: rather, the consumer will wish to wait for someone else to purchase

---

Holder: again, no bargaining will take place, but this time preservation will not occur, because the Rights Holder has the legal right not to preserve, and there are no net benefits the Beneficiary can allocate to the Rights Holder to induce preservation.

[41] Waters, D. (2002) "Good Archives Make Good Scholars: Reflections on Recent Steps Toward the Archiving of Digital Information", in *The State of Digital Preservation: An International Perspective*, p. 78-95. Available at: http://www.clir.org/pubs/abstract/pub107abst.html

[42] Cooter, R., and T. Ulen (1988) *Law and Economics* (Scott, Foresman, and Company, Glenview), p. 46

[43] Of course, the tax-evading citizen may have to enjoy these benefits from the confines of a jail cell!

the good, and then enjoy the benefits without payment. This is the notorious "free-rider" problem. From the supplier's perspective, the cost of excluding non-paying consumers from benefiting from the good is prohibitively high; therefore, the supplier will have little incentive to supply the good.

The free-rider problem impacts digital preservation when the digital materials in question are non-unique: in other words, copies are owned by multiple institutions. This leads to a diminishment in the incentive to purchase (i.e., demand) digital preservation services. Suppose copies of a particular digital resource are owned by many institutions. If one of these institutions chooses to go to the trouble and expense of preserving the resource, the others can "free-ride" on the benefits – once the resource is preserved for one institution, it is perceived as being preserved for all, since the preserved copy is substitutable for all other copies. As Waters observes in his paper, these circumstances lead to a situation where the incentive to preserve is diminished for all owners of the digital resource in question. Since the free-rider problem manifests itself through the collective belief that benefits can be obtained by all parties from an economic activity undertaken (and paid for) by one party, no one wishes to be the "first mover" – the party that undertakes the activity.

The public good qualities of preservation, with its associated impact on incentives, is not unique to the preservation of digital materials, but extends to preservation in any form. However, the public good issue, and in particular the free-rider problem, has the potential to be especially pervasive in the digital realm. That is because there is a large class of digital resources that are owned by multiple institutions, are non-unique (each owner possesses at least one distinct physical copy), and whose preservation confers benefits on most if not all of the owners: specifically, rendering applications and operating systems.

Digital resources can be divided into two broad classes – first, the resources that are actually utilized by the end user (books, music, movies, data sets, etc.), and second, resources that constitute the software environments needed to render and utilize the resources in the first category. From an institutional perspective, much of the content in the first category is either licensed or subscribed, and accessed over a network: therefore, ownership, in terms of both intellectual property rights and physical possession of the bits, is retained by the Rights Holder. In these circumstances, the positive externality issue discussed above is the economic problem of central importance. In contrast, most institutions own non-exclusive rights to, and possess physical copies of, the operating systems and applications needed to render, interpret, and utilize digital resources. Operating systems and applications often take the form of shrink-wrapped software, purchased outright by the institution and delivered on a physical medium such as CD-ROM.

Most owners of copies of operating systems and rendering applications would benefit from the assurance that at least one copy survived over the long-term. Indeed, some digital preservation strategies, such as emulation, require that the software environment within which a digital resource is accessed be preserved along with the

resource itself. But only one copy of a particular software environment, or at most, a limited number of copies, needs to be preserved to ensure that it is available for all. In this sense, there is a strong incentive for each stakeholder to let someone else assume the responsibility and cost of seeing that preservation is carried out.

One might conjecture that the free-riding problem can be easily overcome by stipulating that no one may access the preserved object but those who pay a fee for the privilege. This measure may fail to supply the appropriate incentive, however, unless the benefits from *preservation* are combined (and inseparable) from the benefits of *access*. If institutions can continue to utilize their current, unpreserved copy of the resource for access purposes, rather than relying on access to a preserved copy residing in an archive, they will perceive an interest in continuing the waiting game, in the hope that one of the other institutions will eventually shoulder the burden of taking the necessary steps to preserve the resource. The other institutions can then pay for access to the preserved object later, if needed; in any event, the safety net for their own copy is established. This problem becomes especially cumbersome if there is a great deal of uncertainty over whether or not the viability of current copies are at risk in the short term. For example, if the resource is encoded in a widely-used, non-proprietary format, there may be little concern that the means to render the resource will become unavailable any time soon. This enhances the incentive to wait for another institution to carry out preservation, since immediate action to preserve the object is probably unnecessary.

Why does the free-rider problem not hold for unique digital resources? By definition, a unique resource exists as a single physical copy, owned and controlled by a single Rights Holder. This Rights Holder serves as the sole decision-maker in determining whether or not preservation is carried out, based on their perception of the associated costs and benefits. If preservation is not carried out, the source of market failure is not the free-rider problem – since other institutions with an interest in preserving the resource do not own it, they have no basis for initiating and carrying out preservation unless that right is ceded to them by the Rights Holder. In contrast, the concept of free-riding is based on the idea of multiple entities that *could* take action, but choose not to.

The effects of the free-rider problem on the prospects for the development of an aftermarket in digital preservation services are obvious. Since institutions have little incentive to be the first to take action to preserve a digital object that is simultaneously owned by other institutions, there will be little demand for an aftermarket preservation service, since the non-rivalrous nature of digital preservation suggests that once one institution has preserved the object, all others who own it will also benefit. Each institution will wait for another to make the first move to preserve, and the net effect will be a sharply reduced demand for aftermarket services in the preservation of digital materials.

The free-riding problem can be represented in another way. It may be that in the short- and medium-term, each Rights Holder/Beneficiary entity will find private motivation sufficient to take steps to preserve their own copy of the digital resource in

question, regardless of the actions taken by others. But what if the goal of preservation is extended to include benefits from a societal perspective – for example, to ensure that at least one viable copy of the resource survives for future generations of users? Rather than combining preservation with relatively continuous access, this would be something along the lines of a "time capsule" approach to preservation. In scenarios where the objectives of preservation apply only to the very long-term, preservation assumes the qualities of a pure public good, with the attendant problem of free-riding. Preservation in this form is non-excludable: in preserving one copy of a resource for future generations as part of the scholarly or cultural record, it is impractical to exclude the descendants of current decision-makers who choose not to contribute to preservation from participating in its benefits. Also, enjoyment of the benefits of preservation (i.e., knowing a copy is safely preserved) by one party does not detract from another party's enjoyment of these same benefits. In these circumstances, preservation takes the form of a public good, and no privately motivated entity would be willing to fund it.

In sum, digital preservation can take on properties similar to a public good when distinct copies of the digital resource to be preserved are owned by multiple institutions, each of which has a stake in the preservation of the resource. Therefore, the free-riding problem will tend to arise in scenarios where 1) the roles of Beneficiary and Rights Holder are held by the same entity, and 2) there are multiple institutions filling these dual roles. Referring back to the five organizational models, these conditions have the potential to hold in two models: the Centripetal Model and the Demand-Side Model.

Economic theory suggests a simple solution to the free-rider dilemma. Public goods, such as digital preservation, exhibit two key characteristics: non-rivalry in consumption, and non-excludability. These characteristics combine to yield the free-rider problem. However, by eliminating one of these characteristics, the free-rider problem can be averted. In particular, measures can be adopted to transform public goods into what economists call *club goods*: a non-rival, *excludable* good. A good example of a club good is pay-per-view television. The service is non-rival, since adding one more viewer does not reduce the consumption of the service by other viewers. However, the service is excludable, since those who do not subscribe to the service do not consume it. As such, the market failure caused by free-riding does not impede the development of a viable market for fee-based television movie services.

In the context of digital preservation, the obvious way to prevent Beneficiaries from enjoying the benefits of preservation without incurring a portion of the cost is by denying access to the archived materials to those who do not contribute to preservation. But as noted above, exclusion based on access is difficult to achieve if each Beneficiary owns and controls its own copy of the resource. In this case, each Beneficiary can rely on its own copy for access, while waiting for one of the other Beneficiaries to take the initiative to preserve. The free-riding problem would then persist.

To overcome this problem, preservation must somehow be made inseparable from access. This can be achieved by altering what is meant by purchasing or owning a digital resource. In particular, some publishers of digital content allow Beneficiaries to purchase

the "right to access", rather than an actual physical copy of the resource. While the Beneficiary can freely access the copy through the publisher's server, it is not permitted to mount a separate copy on its local system. In this scenario, it is then simple to overcome the free-riding problem associated with multiple owners of the same digital resource. Beneficiaries cannot rely on their own physical copy of the resource for access, since they do not possess such a copy: all Beneficiaries access what is essentially the same copy of the resource through the publisher's central server. If a particular Beneficiary "opts out" of the cost of preservation by refusing to pay fees, they also opt out of access to the resource, since in this case, preservation and access are inextricably intertwined. The service terms of JSTOR are an excellent example of the "club good" approach. Libraries and other institutions can purchase access to content in JSTOR's archive, but this content cannot be mounted locally: it must be accessed through JSTOR's server. Access and preservation fees are bound together: both must be paid to obtain access to the content. By altering the conditions of ownership in order to make preservation and access inseparable, the free-riding problem is effectively avoided, despite the fact there are multiple entities who would benefit from preservation.

But the JSTOR model cannot be applied in all digital preservation scenarios. "Time-capsule" preservation – where the object of preservation is not to ensure continued access on the part of the current generation of users, but rather, to ensure at least one copy of the resource survives for future generations – is a textbook case for the free-rider problem. If many institutions own copies of the resource, an initiative by any one of them to preserve it for future generations would satisfy the objectives of preservation. But there is little incentive for any institution to make the move to do so.

Economic theory suggests remedies for situations where preservation and access are separable. Manifestations of the free-rider problem in these instances can be overcome by government provision of the digital preservation activity, financed through the regular system of compulsory taxation. In this way, the desired preservation objective is met, and the cost is distributed relatively evenly among stakeholders.

V.5.c. Heterogeneous Demand, Spillover Benefits, and Economies of Scale

Digital preservation is an issue impacting a diverse community of stakeholders, among which there is no discernable consensus on what preserving digital materials over the long-term precisely means. Differentiation in the perception of digital preservation's anticipated outcome can arise from a number of sources. For example, suppose two institutions both own distinct copies of a particular digital resource. Furthermore, both are stakeholders in the long-term viability of the resource. However, their respective definitions of "long-term" differs. One institution believes that the resource is likely to be valuable for approximately ten years. The other institution believes the resource will be valuable "in perpetuity".

Alternatively, differentiation can be based on the degree or "intensity" of preservation. To one institution, preservation of a particular digital resource might involve perpetuating the resource's intellectual content over time, in whatever format is convenient and economical; to another institution, preservation of this same resource

might involve preserving not only the resource's intellectual content, but also its original look, feel, and functionality – in other words, an exact replica of the object at the time it was created.

Clearly, the type of digital preservation service offered will have important implications for its long-term cost. Preservation that is carried out for extremely long periods, or that attempts to maintain original features, functionalities, or environments associated with archived resources, will almost certainly be more expensive to carry out than preservation lasting for periods of shorter duration, or involving alteration of the archived resources to suit available technologies or other elements of the contemporary environment.

Differences in expected overall preservation costs will feed back into the behavior of the entities – the Beneficiaries – who constitute the demand for digital preservation services. It is reasonable to assume that those who place higher value on the long-term retention of a given digital resource would be willing to pay for higher levels of digital preservation service. Conversely, those who value long-term preservation less would only wish to pay for lower levels of preservation service.

Consider an entity – an Archive – who wishes to provide digital preservation services to a group of Beneficiaries, some of which demand a high-level preservation service, and others of which demand a low-level preservation service. This heterogeneity in the demand for digital preservation services poses a dilemma for the Archive. Suppose the Archive chooses to offer high-end services (e.g., retention in perpetuity, preservation of original "look and feel"). In this case, the Archive will capture the "high-end" of the market – those who value preservation relatively highly and exhibit a corresponding willingness to pay for more sophisticated forms of preservation. But the "low-end" of the market – those who do not place great weight on preservation, and therefore are willing to purchase only limited preservation services at lower cost – will be priced out of the market. This impacts the Archive in two ways. First, the Archive foregoes the revenues that could have been earned by offering a low-end service to the low-end of the market. Second, offering only the high-end service could severely limit the size of the market, which could in turn have a profound impact on the costs of providing the preservation service. In particular, if the remaining market is not sufficiently large to allow the archive to realize cost efficiencies through economies of scale, the rate of return from offering digital preservation services would be reduced. In the extreme, the Archive may decide that the returns from offering the preservation service at a small scale preclude offering the service at all.[44]

One alternative open to the Archive is to provide the high-end service to the entire market at the low-end price. But in these circumstances, it is unlikely that the Archive would be able to re-coup its production costs.

[44] This would likely depend on the financial objectives governing the Archive's activities (e.g., profit-maximization, cost-recovery).

Another option is for the Archive to offer the low-end service (retention for a period of finite duration, preservation of intellectual content only), rather than the high-end service. This would induce the entire market to participate: customers for low-end services would be willing to pay the lower cost for limited preservation services, and customers for high-end services would purchase the low-end services because limited forms of preservation are better than no preservation at all.[45] But although the market may now be of sufficient size for the Archive to achieve economies of scale, it would forego the opportunity to earn greater revenues from customers at the high-end of the market, who are willing to pay more for a higher level of preservation service. Furthermore, the preservation needs and requirements of the high-end of the market would not be fully met. The latter consequence may be of especial importance to mission-driven Archives created to serve the cultural heritage or scholarly communities.

The difficulty in serving a market of multiple Beneficiaries with heterogeneous perceptions of what digital preservation entails can significantly reduce the attractiveness of entering the aftermarket for digital preservation services to would-be Archives. Put another way, the incentives to preserve, in the sense of the willingness to provide digital preservation services, could be diminished.

This incentive problem arises when the Archive role and the Beneficiary role are held by separate entities. If the Archive and the Beneficiary are the same entity, the problem of heterogeneous demand does not arise, since a single entity will both articulate the desired level of preservation service, and carry it out. The Archive serves one customer: itself; therefore, a single level of preservation service will prevail. But when the Archive is a distinct entity from the Beneficiary, there are multiple Beneficiaries, and there is a varying perception among these Beneficiaries as to the level of preservation service needed for digital resources, the incentives for the Archive to provide digital preservation services may be reduced or eliminated. These circumstances can occur in the Centrifugal Model, the Supply-Side Model, and the Demand-Side Model.

Economic theory suggests the following strategy for addressing this incentive problem: offer different versions of the service to different segments of the market. "The key," observes Shapiro and Varian, "is to identify dimensions of your … product … that are highly valued by some customers yet of little importance to others. Then offer versions that differ noticeably in ways designed to appeal selectively to each type of customer."[46]

In the context of digital preservation, the Archive could offer preservation services for time periods of varying lengths, including perpetuity. It could also offer a range of preservation services, from simple bit storage, to preservation of intellectual content, to preservation of the original resource's precise "look and feel." The idea is to differentiate the service menu in such a way that customers self-select based on their particular preservation needs.

[45] Assuming, of course, there are no competing Archives that are willing to offer the high-end service to a limited market.
[46] Shapiro, C., and H. Varian (1999), p. 55-56

Two problems emerge, however, when applying the strategy of market segmentation to digital preservation. First, implicit in this strategy is the assumption that the service provider can exclude low-end buyers from enjoying spillover benefits from the provision of the high-end service – a situation reminiscent of the free-rider problem discussed in the previous section. Suppose two institutions want to preserve their respective copies of the same digital resource. One institution would be satisfied with the low-end service (e.g., preservation for 10 years); the other requires the high-end service (preservation "in perpetuity"). To obtain their respective service needs, both institutions pay an upfront, capitalized archiving fee. Presumably, the cost to the first institution will be less than that of the second.

The net result of the two preservation agreements will be that the archive preserves the resource forever, since the terms of the second agreement subsume the terms of the first. In this event, the first institution has knowingly or unknowingly "hedged its bet" that the resource will no longer be worth preserving after 10 years. Should this prediction prove wrong, the first institution has a safety net, since the resource has been preserved beyond the ten-year horizon due to the wishes of the second institution. This is achieved without committing to the higher upfront archiving fee associated with preserving the resource in perpetuity.

If it becomes common knowledge that at least one institution has taken steps to preserve a particular digital resource using a high-level service, then all Beneficiaries may rest assured that the resource will be fully preserved, regardless of their own preferences on the appropriate level of preservation for the resource. These Beneficiaries will then have a strong incentive to purchase more basic – and less expensive – forms of preservation for their own copies of this same resource. Should circumstances turn out that their preservation needs exceed the features offered by the low-end service, the earlier decision to purchase only low-end digital preservation services is reversible: a copy of the resource exists that has been preserved according to the parameters of the high-end service. The consequence will be a diminished demand for the high-level preservation service. We have then returned to the free-rider problem of the previous section.

Ideally, the Archive would like to exclude institutions from enjoying this unpaid-for benefit. But to do so is difficult. A digital resource cannot be both preserved only for a finite period, *and* in perpetuity. If the Archive is obligated to preserve a digital resource for a long period to meet agreements with one set of Beneficiaries, it unintentionally confers an unpaid-for benefit on others who are only willing to preserve the resource for some shorter period at lesser cost. This suggests that the Archive will have difficulty motivating Beneficiaries to reveal their true preferences for long-term preservation, and by extension, to pay the higher upfront capitalized archiving fee. To counteract this incentive problem, segmentation of the market into varying levels of digital preservation service must be accompanied by *incentive compatibility constraints*. The object of incentive compatibility constraints is to ensure that in a market characterized by differentiated demand, consumers choose a version of the product commensurate with

their true preferences. In the case of digital preservation services, this might take the form of a fee levied in the event that a Beneficiary wishes to gain access to a copy of a digital resource preserved at a higher level than the service actually purchased.[47] This fee must be sufficiently high such that the Beneficiary finds the combined cost of low-end preservation and future access fees more burdensome than purchasing the high-level service upfront.[48]

In addition to the problem of excluding low-end buyers from enjoying spillover benefits from the high-end service, another issue arises: can different levels of digital preservation service be provided by a single Archive in an efficient way? This question extends from digital preservation's public good attributes – non-rivalry in consumption and non-excludability. Typically, goods with these characteristics are most efficiently supplied in a "one-size-fits-all" version. Would it be efficient, for example, to offer different service levels for national defense, paid for by different levels of taxation? In the same way, would it be efficient for an archive to preserve multiple copies of a digital resource, distinguished on the basis of varying intensities of preservation? It may be that if the market for digital preservation services does move toward segmentation based on the level of service, production efficiency could be achieved through specialization: archives would offer only one level of archiving service, rather than attempting to offer a menu of different levels of service. As more is understood about the technical processes associated with various levels of digital preservation service, further research is needed to identify potential economies of scale within these processes that would tend to favor a specialization approach.

V.6. Prospects for Remedies

In considering the prospects for devising appropriate remedies for the potential market failures, or incentive gaps, associated with digital preservation, it is useful to keep in mind that the community is currently operating on the cusp: on the one hand, making arrangements to provide retrospectively for the long-term preservation of existing digital materials, while on the other, evaluating means to formalize this process for future materials. These efforts are occurring in what is for the most part an absence of a well-defined market for digital preservation services. It is difficult at this point, therefore, to make definitive statements bearing on the nature and magnitude of the remedies necessary to ensure an appropriate level of digital preservation activity.

It is interesting nevertheless to consider real-world digital preservation activities in light of the incentives to preserve, potential market failures, and, where needed, remedies to enhance incentives. This section briefly discusses two very different digital

[47] Who actually collects this fee depends on the circumstances. If it is the Beneficiary who preserved the resource using a high-level service, an interesting possibility emerges: one can imagine institutions incurring the cost of expensive, high-level preservation as something akin to an investment, anticipating that other institutions will eventually require, and be willing to pay for, access to this "intensely preserved" resource.

[48] Intellectual property rights can also serve as an incentive compatibility constraint. If copyright restrictions prevent one owner of a digital resource from purchasing access to another owner's copy (i.e., the "first sale doctrine" is not applicable), then the fact that the second owner purchased a higher level preservation service will not benefit the first owner.

preservation activities, relating them to the organizational models of Section IV, noting potential sources of market failure inherent within these activities based on their organizational structure, and discussing how sufficient incentives to preserve were achieved and maintained.

V.6.a. Libraries, Publishers, and JSTOR

JSTOR[49] is a non-profit organization that "acquires rights from publishers to full runs of selected journals in the humanities and social sciences, digitizes the content, and makes it available on the Web through institutional site licenses."[50] JSTOR is both an access and an archiving service: subscriptions to JSTOR's collections include annual access fees, as well as a one-time, capitalized archiving fee. Although JSTOR provides access to the digitized content, publishers retain the rights to the content of their journals. In addition, JSTOR users (e.g., library patrons) have only limited access privileges: in particular, access to the content in JSTOR is confined to personal and non-commercial use, and users are prohibited from downloading "an entire issue of a journal, significant portions of the entire run of a journal, a significant number of sequential articles, or multiple copies of articles."[51]

In the JSTOR environment, the three decision-making roles are embodied in three distinct entities: JSTOR itself as the Archive; libraries and their patrons as the Beneficiaries; and commercial publishers as the Rights Holders. This suggests that the JSTOR activity is represented by the Centrifugal Model. As noted above, the Centrifugal Model harbors the potential for two sources of market failure: positive externalities and variation in demand. How did JSTOR overcome these pitfalls?

The first factor serving to diminish the incentives to preserve is the fact that in this scenario, the Rights Holder and Beneficiary roles are held by distinct entities, leading to the potential for a positive externality. To publishers, the benefits of funding long-term preservation of their content (perhaps in the form of cultivating goodwill among subscribers) are likely outweighed by the substantial costs involved. Libraries, on the other hand, see greater benefits, stemming from their mission to provide ongoing access to information. But these benefits are not internalized by the publishers. What then is the incentive for publishers to allocate resources to preservation?

This apparent lack of sufficient incentives was overcome by the creation of what is in effect a subsidy from Beneficiaries to the Rights Holder: in particular, Beneficiaries, through the payment of access fees and the one-time archiving fee, subsidize completely the cost of preserving the journal content (in the form of digital surrogates), thus relieving publishers of the trouble and expense. Of course, the subsidy is not paid directly to publishers: rather, it is paid to JSTOR, the entity responsible for carrying out the preservation process on behalf of both Beneficiaries and Rights Holders. All parties benefit from this arrangement. Beneficiaries gain assurance that the journal content to which they subscribe will be preserved and made available for long-term access: even if a

[49] http://www.jstor.org/
[50] "Building and Sustaining Digital Collections: Models for Libraries and Museums", p. 3-4
[51] JSTOR Terms & Conditions of Use: http://www.jstor.org/about/terms.html

publisher later chooses to terminate their agreement with JSTOR, the latter may continue to provide access to the content archived up to that date to all institutions who purchased access prior to termination of the agreement. Publishers benefit from the fact that they can preserve their content at no expense to themselves, retain the rights to their journal content, protect existing revenue streams[52], and cultivate goodwill among institutional subscribers. Finally, JSTOR can fulfill its stated mission, which is to "build a reliable and comprehensive archive of important scholarly journal literature."[53]

A second factor potentially working to diminish incentives to preserve in the JSTOR environment is the fact that the roles of Beneficiary and Archive are held by distinct entities, and there are multiple Beneficiaries. This leads to the possibility that variations in the nature of the demand for preservation services among subscribers could fragment the market to the extent that JSTOR is unable to achieve cost efficiencies through economies of scale. To counteract this possibility, JSTOR offers what is essentially a one-size-fits-all service to all subscribers. But this approach is largely explained by the nature of the digital materials preserved by JSTOR. JSTOR digitizes paper-based journals: therefore, the digital objects are static representations of text-based documents. There is little question of preserving the original "look and feel" of the source documents, since this is, for all intents and purposes, captured in the image file. Given this, it is unlikely that a demand for varying levels of service along these lines would exist among Beneficiaries. In terms of the period of archival retention, there is more scope for variation in demand in this regard. But in practice, all subscribers are required to support what is essentially preservation "in perpetuity" through payment of the one-time archiving fee, in addition to an annual access fee. The required archiving fee is enforced through the access functionality: no access to JSTOR collections is permitted without payment of the archiving fee. It would be interesting to determine whether this required archiving fee has significantly reduced demand for JSTOR services. For example, would additional institutions be motivated to subscribe to JSTOR if a "pay-per-article" scheme, with no archiving fee, were offered? Would other institutions still be willing to pay the archiving fee to ensure the long-term preservation of JSTOR's collections? How could those who choose not to pay the archiving fee be excluded from the benefits produced by those who do?

V.6.b. Outsourcing Preservation: the JPMorganChase I-Vault! Service

JPMorganChase, a financial services firm, operates I-Vault![54], a service that will store and maintain imaged documents on behalf of their owners for extended periods of time. Examples of materials that might be deposited into I-Vault! include, but are not limited to, proprietary corporate documents, bank records, student transcripts, and medical records.

The types of materials the I-Vault! service is positioned to accommodate are generally proprietary and confidential in nature. Since it is likely that access to these

[52] JSTOR's collections are governed by a "moving wall" between the last issue archived and the current issue; this moving wall is generally between 3 to 5 years.
[53] JSTOR's Mission & Goals: http://www.jstor.org/about/mission.html
[54] http://www.jpmorgan.com/cm/cs?pagename=Chase/Href&urlname=jpmorgan/ivault

materials would not be permitted beyond the Rights Holder, it follows that the Rights Holder is also the Beneficiary – indeed, the sole Beneficiary – from their preservation. The Archive role is, of course, held by I-Vault!, which is advertised as an outsourcing solution, and is therefore an entity distinct from the Rights Holder/Beneficiary. The I-Vault! environment reflects the Demand-Side Model: i.e., a scenario where the roles of Beneficiary and Rights Holder are combined into one entity, while the role of Archive is held by a separate entity[55].

The Demand-Side Model harbors two sources that could potentially diminish the incentives to preserve: the "free-rider problem", and fragmented demand that inhibits the realization of economies of scale. How has the I-Vault! service fared in regard to these issues?

In terms of the free-rider problem, the answer is evident from the types of materials I-Vault! is intended to house: specifically, proprietary, and probably confidential materials owned and controlled by a single Rights Holder. This Rights Holder is also a Beneficiary in the preservation of the digital materials in question, and while there may be other Beneficiaries as well, the Rights Holder is the only Beneficiary that possesses the right to initiate preservation. In the extreme, the Rights Holder may be the sole Beneficiary, as would be the case with confidential corporate records and similar materials. The free-rider problem is avoided by the absence of *multiple* entities combining the Rights Holder and Beneficiary roles in regard to a particular digital resource. In this sense, the incentive to wait for another Rights Holder to initiate preservation is removed.

In the I-Vault! environment, the Archive and the Beneficiary roles are held by separate entities. This suggests that if Beneficiaries exhibit heterogeneity in the type of digital preservation services they require and are willing to pay for, the Archive may not be able to realize necessary economies of scale in the provision of one or more service levels. The I-Vault! service does in fact offer some variation in its service plans: for example, the I-Vault! Web site notes that "[r]etention period, storage media, and index (search) criteria can be customized to meet your specific requirements."[56] Did offering the opportunity for customization fragment I-Vault!'s market in such a way that archiving costs, and even the demand for archiving services, were adversely impacted?

Casual examination of the I-Vault! environment suggests that this was not the case. As discussed earlier in the paper, one problem associated with the provision of a range of preservation service levels is that the free-rider problem crops up again: in cases where multiple entities own the same resource, each has the incentive to let another entity pay for more expensive (and presumably more sophisticated) levels of preservation. This

[55] The characterization of I-Vault! as an entity distinct from the Rights Holder and Beneficiary is not quite accurate, because one of its largest customers is in fact JPMorganChase. In that context, I-Vault! would reflect the Centripetal Model where all roles are combined into one entity. However, I-Vault! is also offered as an outsourcing solution for external customers, and it is in this context that the discussion above is placed.

[56] http://www.jpmorgan.com/cm/cs?pagename=Chase/Href&urlname=jpmorgan/ivault/techinfo

may reduce demand for particular levels of preservation service, or even overall demand for preservation services. However, the types of materials typically deposited into I-Vault! (i.e., proprietary and confidential) seem to preclude this possibility. Since materials archived in I-Vault! are likely to have no cross-institutional overlap, there is likely no need to archive the same resource according to multiple service customizations.

What about achieving economies of scale? Several factors seem to work in I-Vault!'s favor in this regard as well. While certain aspects of the service are customizable, the basic infrastructure – i.e., ingest procedures, storage systems, access mechanisms, etc. – seems to support all service customizations. It is also key that I-Vault! handles essentially one class of digital object – digital images – which serves to reduce the scope of variation in demand on the part of organizations utilizing its service. Because the archived objects are digital versions of static documents, the strategy for securing their preservation is largely uniform and straightforward – there is no "functionality", for example, that must be maintained in working order over the long-term. This simplifies the preservation process, in that there is no need to preserve a portion of the archived objects using a strategy such as migration to preserve intellectual content only, while preserving the remainder through a different preservation strategy such as emulation to preserve look, feel, and functionality. It is likely that these factors contribute toward I-Vault!'s ability to operate at very high volumes: according to its Web site, I-Vault! currently archives the equivalent of over 700 million document pages.

* * *

In considering the interplay between organizational models for digital preservation activities, and remedies to correct for insufficient preservation incentives, the following question can be posed: is the remedy a consequence of the organizational model, or *vice versa*? The answer is likely "both".

In some circumstances, the prevailing organizational model may be a by-product of negotiations among stakeholders on how a particular digital preservation activity is to be carried out. The recent agreement between Elsevier and the National Library of the Netherlands is a good example.[57] The Library will serve as a digital archive for copies of all Elsevier journals made available through the ScienceDirect service. The purpose of this agreement is to ensure the preservation of the scientific heritage for the public; the Library serves as the representative of these stakeholders, and also as the archive itself. The organization of this initiative seems to reflect the Consolidated Model, where the roles of the Archive and Beneficiary are held by one entity[58], and that of the Rights Holder by another. Since the roles of Rights Holder and Beneficiary are manifested in separate entities, incentives to preserve are diminished by the presence of a positive externality. But this preservation activity has a built-in remedy, in the form of a subsidy from the combined Archive/Beneficiary entity to the Rights Holder to cover the cost of

[57] For a description of this partnership, see:
http://www.kb.nl/kb/resources/frameset_kb.html?/kb/pr/pers/pers2002/elsevier-en.html

[58] We are assuming here that the National Library of the Netherlands "represents" a community of Beneficiaries, and acts on their behalf.

preservation. Specifically, the National Library of the Netherlands will assume the responsibility (and presumably the costs) of carrying out preservation. Relieved of this burden, Elsevier's incentives to participate in the preservation process are enhanced: from Elsevier's perspective, the benefits, in the form of cultivating goodwill among its customers, now likely outweigh the costs, which are substantially reduced. In this sense, the terms of the agreement concluded between Elsevier and the Library motivated the organizational model, but also specified a means to create appropriate incentives within that model.

In other scenarios, however, the organizational model may be imposed rather than adopted. For example, consider digital materials obtained through license or subscription. It may be that the Rights Holder refuses to cede the right to preserve the materials to an independent Archive, or a combined Archive/Beneficiary entity. In this event, any digital preservation activity that eventually emerges will be conducted under the auspices of the Supply-Side Model, where the Beneficiary role is held by a distinct entity, and the Rights Holder and Archive roles are combined in another. Remedies which are subsequently devised to correct perceived insufficiencies in the incentives to preserve must work within the organizational model bestowed by circumstances, rather than serving as a means of shifting the preservation activity to another organizational model.

An important topic for future research should be to examine how the organizational structures of digital preservation activities come about, with an emphasis on the factors which tend to motivate the adoption of one model over another. A starting point for this research, and one that would contribute significantly toward a proper allocation of incentives among relevant decision-makers, would be a formalization of the structure of, and responsibilities for, digital preservation activities in various stakeholder communities. As it becomes clearer which of the organizational models will prevail in a particular community, and within the framework of this model, which of the key decision-makers will take the lead in securing the long-term preservation of digital materials important to the community, it will be correspondingly easier to determine whether sufficient incentives to preserve naturally emerge from this scenario, and if not, what remedies are appropriate to address the problem. This in turn will inform the development of a robust market for digital preservation services that can efficiently serve the community's needs and requirements.

VI. SUMMARY, AREAS FOR FUTURE RESEARCH, AND CONCLUSION

Digital preservation tops the agendas of many institutions in the cultural heritage and information management communities. As the number of digital assets which they own outright, or for which they are at least stakeholders, continues to expand, the need to take measures to secure their long-term availability grows commensurately. Technical challenges remain in regard to achieving this objective. But as digital preservation moves beyond the realm of small-scale, experimental projects to become a routine component of a digital asset's life-cycle management, the question of how it can be shaped into an economically sustainable process begins to overshadow other concerns.

Digital preservation can hardly be classified as a new topic anymore, yet we still find ourselves not very far from the beginning in terms of exploring its economic ramifications. No systematic study of the economics of digital preservation has yet emerged. While some progress can be reported in terms of developing cost models, and to a lesser degree, strategies for recouping costs, there has been no general research into how digital preservation might translate into a viable economic service, demanded and supplied through ordinary market mechanisms. The lack of research on this topic has persisted even as digital archiving services begin to appear.

This paper is intended to serve as a foundation for future research in this area, by focusing on the most fundamental economic question: *what are the incentives to preserve digital materials?* It begins by identifying three key decision-makers associated with the process of digital preservation. Next, these decision-makers are combined and recombined into a number of simple organizational models serving as (extremely) broad descriptions of the various scenarios under which digital preservation activity might take place. With these models in hand, the paper moves on to examine the prospects of providing digital preservation as a market-based service. The approach taken here was to consider digital preservation as an aftermarket service – in other words, a service that enhances the value and longevity of a digital information object. Using this perspective as a benchmark, economic intuition can be applied to evaluate whether, in context of the five organizational models, sufficient incentives exist on the part of relevant decision-makers to nurture the development of an aftermarket in digital preservation services.

A sobering conclusion from the analysis is that the various organizational models often give rise to insufficient incentives to undertake digital preservation. Even without reference to quantified cost and revenue models, intuition based on economic theory suggests that in many circumstances, there is a strong likelihood that the key decision-makers in digital preservation simply have little or no motivation to participate in the process. This in turn suggests that the development of a functioning market in digital preservation services would be expedited by the development of policies aimed at matching preservation objectives with incentives.

But to do so requires a great deal of additional (and more focused) research. What follows is a proposed agenda for future research into the economics of digital preservation.

VI.1. A Four-Point Agenda for Research into the Economics of Digital Preservation

1) Accumulation and Synthesis of Real-World Case Studies in Digital Preservation

A great deal of insight, as well as both quantitative and qualitative data, can be obtained through the accumulation of case studies of real-world digital preservation initiatives. These case studies should represent a broad cross-section of digital preservation activities, differentiated on the basis of scale, scope, technology, and other criteria, and may also extend to other relevant industries, such as data warehousing.

Objectives include:

- Analysis and comparison of alternative preservation strategies

- Analysis and comparison of alternative structures for digital preservation activities
- Syntheses that highlight important points of convergence and divergence among the case studies
- Specification of cost and revenue models
- Development of metrics for consistent cross-study analysis
- Identification of points of convergence and divergence between the preservation of digital materials on the one hand, and analog materials on the other.

2) Incentive-Boosting Policies for Digital Preservation

The discussion in this paper indicates that in many scenarios, insufficient incentives exist to carry out digital preservation. Research is needed to develop policies and programs aimed at enhancing the incentives to preserve in situations where sufficient incentives do not emerge naturally.

Objectives include:

- Understanding of the relationships between various rights management policies and the incentives to preserve
- Matching corrective policy measures with the types of market failures for which they are the most effective
- Demarcation of the appropriate scope for non-market provision of digital preservation services (i.e., through public agencies), rather than through incentive-based market mechanisms
- Understanding of how the organizational structure of digital preservation activities comes about, and also how it evolves over time.

3) Market Structure for Digital Preservation

Overcoming the incentive problems associated with the development of an aftermarket for digital preservation services leads to the question of the form such a market might ultimately take. Research should be aimed at understanding the structure and characteristics of an efficient global market in digital preservation services.

Objectives include:

- Understanding the incentives to centralize: what is the scope for creating centralized services for carrying out key parts of the digital preservation workflow, or more generally, the digital asset management life cycle? Identify presence of fixed costs that could serve as sources of economies of scale.
- Understanding the incentives to specialize: is digital preservation more efficiently carried out as an integrated process carried out by a single entity, or as a package of specialized services provided by multiple entities?
- Should digital preservation be a "stand-alone" service, or should it be integrated with other digital asset management services? Would such an integration serve to lower costs?

4) Pricing Strategies for Digital Preservation Services

Economically sustainable digital preservation services require effective strategies for recouping the costs of preservation. Research must be aimed not only at developing pricing mechanisms for digital preservation services, but also at specifying the exact characteristics of the services to which the pricing mechanisms apply.

Objectives include:

- Demand studies/market surveys that articulate demand-side expectations for digital preservation services: for what are customers willing to pay?
- Specification of various levels of digital preservation services, including characterizing the anticipated outcomes of these services and the nature of the service "guarantees" to which service providers can realistically commit
- Establishing "minimum digital preservation service requirements" to meet the needs of particular communities
- Devise pricing mechanisms that accurately reflect the true cost of managing the long-term retention of digital materials
- Devise pricing mechanisms that create an appropriate distribution of risk between the owners of digital materials and the digital preservation service provider

VI.2. Conclusion

Preservation of analog materials has traditionally been achieved through something akin to a "crisis management" strategy. Once conditions have reached a point where there is significant risk that valuable materials will be lost, efforts are organized to stave off impending disaster. These efforts are often financed through sporadic infusions of "soft", unbudgeted funding – government grants, philanthropic donations, etc.[59] The Brittle Books Program established in the 1980s by the US National Endowment for the Humanities is illustrative of this approach.

It is likely that one-off funding mechanisms will also play a prominent role in supplying resources for digital preservation. But the characteristics of digital materials are such that primary reliance on an *ad hoc* approach to the economics of digital preservation is almost certain to prove inadequate. The time horizon beyond which preservation becomes a concern can be extremely short – indeed, one could argue that the preservation process for digital materials begins at the moment of creation. It may be difficult, therefore, to separate the preservation process from those associated with the short-term, or even day-to-day, management of digital materials. In short, it is likely that digital preservation programs will assume a much more activist character than those dealing with analog materials. To meet this contingency, economic mechanisms need to be in place to ensure the regular availability of funds, over long time frames, to support preservation processes that are part of a digital resource's routine lifecycle management.

[59] Budgeted funds are also provided for preservation, covering such activities as binding, environmental control, and repair. For libraries, however, these funds typically constitute an extremely small proportion of the institution's overall budget.

As digital preservation continues to evolve from research topic to practical concern, it is likely that the demand for digital preservation services will grow accordingly. But in order for a market in these services to develop and expand, the appropriate incentives to motivate supply and demand must be in place. In circumstances where the characteristics of digital materials and/or the structure of digital preservation activity work to diminish these incentives, measures must be devised and implemented to supply the necessary motivations. Ensuring that appropriate incentives to preserve exist is the fundamental economic imperative of digital preservation, and an essential first step in securing the future of the vast corpus of information in digital form.

* * *

Chesterfield County, in the state of Virginia, is the place where the Bermuda Hundred battlefield sites may be found. A visit to the county Web site reveals a list of "Frequently Asked Questions" about historic preservation under the county's preservation zoning process.[60] The last question on the list is:

What incentives are available to encourage me to restore my property?

In considering the prospects for the long-term preservation of valuable digital materials, it is well to begin by posing a similar question within the digital preservation community.

[60] http://www.co.chesterfield.va.us/CommunityDevelopment/Planning/lndmrk.asp