

94806 A1 – PRIVACY IN THE DIGITAL AGE (Remote course)

Course packet

Content

- Syllabus (this document, on Canvas)
- Readings (on Canvas)

[This page intentionally left blank]

CARNEGIE MELLON UNIVERSITY
HEINZ COLLEGE

94806 A1 – PRIVACY IN THE DIGITAL AGE

Fall 2020 – Mini 1

Monday and Wednesday, 8:00AM-9:20AM, on Zoom

Syllabus

Instructor:

Alessandro Acquisti

acquisti@andrew.cmu.edu

412-268-9853

Office Hours: On zoom (see information on Canvas). Monday 5:00PM-6:00PM and
by appointment

<http://www.heinz.cmu.edu/~acquisti/>

TAs:

Eduardo Abraham Schnadower Mustri

eschnado@andrew.cmu.edu

Office hours: See information on Canvas

Saharsh Agarwal

saharshagarwal@cmu.edu

Office hours: See information on Canvas

TEXTBOOKS

We will not use a textbook in this class, but we will have plenty of academic readings (articles) – with publication dates spanning from 1890 to 2020. All readings have been made available on Canvas.

COURSE DESCRIPTION

Privacy is a complex and multi-faceted concept. This course combines technical, economic, legal, psychological, ethical, and policy perspectives to present a **holistic view** of its role and function in the digital age.

Advances in technology have led to concerns about the reasonable expectation of privacy since Warren and Brandeis wrote their seminal article on the subject at the end of the 19th century. These concerns have evolved and amplified as the technology of communication has been seen as a technology of surveillance. The modern world of computers, cell phones, surveillance cameras, and the Internet of Things offers unprecedented opportunities for tracking everything everyone does.

At the same time, policy around the right to privacy and indeed the definition of what privacy means have evolved in different ways in different countries. European laws attempt to protect the privacy of the individual from corporations, while U.S. law tries to protect the privacy of the individual from the government.¹ Corporations doing business in multiple jurisdictions find themselves subject to conflicting and sometimes contradictory rules and regulations, while users find it difficult to know what rights they have with respect to their interactions.

The reduction of the cost of storing and manipulating information has led organizations to capture increasing amounts of information about individual behavior. New trade-offs have emerged for parties involved with privacy-enhancing or intrusive technologies: individuals want to avoid the misuse of the information they pass along to others, but they also want to share enough information to achieve satisfactory interactions; organizations want to know more about the parties with whom they interact, but they do not want to alienate them with policies deemed as intrusive. Is there a “sweet” spot that satisfies the interests of all parties? Is there a combination of technological solutions, economic incentives, and legal safeguards that is acceptable for the individual and beneficial to society?

This course tries to address the above questions. In particular, the course begins by comparing early definitions of privacy to the current information-focused debate. It then focuses on

- technological aspects of privacy (privacy concerns raised by new IT such as the Internet, wireless communications, and computer matching; tracking

¹ The European Union’s (EU) GDPR Regulation, which came into force May 25, 2018, substantially tightens and toughens requirements on businesses storing, sharing, sending and receiving data of EU citizens. The California Consumer Privacy Act, CCPA, is effective January 1, 2020, also may have a significant impact on corporate privacy initiatives across all sectors of the technology, media and entertainment, and telecommunications industries.

- techniques and data mining; privacy enhancing technologies and anonymous protocols, ...),
- economic aspects (economic models of the market for privacy, financial risks caused by privacy violations, the value of customer information, ...),
 - legal aspects (laissez-faire versus regulated approaches, US versus EU legal safeguards, ...),
 - managerial implications (the emerging role of Chief Privacy Officers, compulsory directives and self-regulative efforts, ...), and
 - policy aspects (trade-offs between individual privacy rights and societal needs, ...)

Let me repeat a key word from the above description: “holistic.” This term refers to the fact that we will try to cover **a variety of angles** in the privacy debate, and we will try to connect and contrast each angle to the others. The approach of this class therefore somewhat privileges breadth over depth (while offering tools and directions for more focused analysis), in order to give you a broad and varied understanding of privacy problems in a networked digital society. (For this reason, this class also provides some preparation and introduction for/to other courses you may take at the Heinz College.)

The goal of this class is not to tell you whether privacy is important or unimportant. Also, the goal of this class is not to scare you into believing that every instance of information disclosure is a privacy invasion – or to convince you that every instance of information collection is legitimate. Instead, the goal of this class is to **educate you about the kinds of information that may be gathered about individuals, and to help you determine the individual and societal trade-offs associated with accepting or avoiding its collection and/or use.** Since privacy is a very subjective and contextual concept, what may be a privacy invasion for you in one setting may not be in another. The challenge we present to you in this course is to understand the difference.

OBJECTIVES

The main objective of this course is to provide **an informed and critical view of the role and value of privacy in the digital age.** Because privacy is a complex and multi-faceted concept, the course aims to present and combine technical, economic, legal, and policy perspectives.

REMOTE COURSE

This entire course will be taught remotely, via Zoom. You will find links to the live Zoom-streamed classes on Canvas. We will also use Canvas for online discussions (via Discussions), for HWs submission, for short quizzes (see below), and for grading. We will strive to make the remote-teaching experience as engaging and interactive as possible. Thus, if you can, please keep your camera ON (so that we can see each other during lectures), and be ready to participate in the class discussion by asking, or answering, questions.

GRADING

The course consists of a combination of readings, assignments, and class discussions. Assignments include homework assignments and a final project.

Specifically, your grade will depend on the following components:

1. Class attendance and short quizzes	10%
2. Homework	40%
3. Project	50%

1. Class attendance and participation include, in addition to your attending Monday and Wednesday classes on Zoom (via Canvas), your participating in class discussions about the readings and topics listed in this syllabus. The readings listed in the schedule of classes should be completed prior to the class for which they are listed. **Short Quizzes** will be administered via Canvas during each lecture, typically close to the end of a Lecture (excluding Lecture 1, of course). They will be very short, and mainly aimed at making sure that you have been able to follow the material covered during the Lecture. Your worst-performing Short Quiz across all lectures will be eliminated (that is, it will not count towards your grade).

2. There will be three homework assignments. Homework assignments will be distributed via Canvas, and should be submitted via Canvas by the deadline indicated on the assignment. Homework assignments consist of short reports (with a predetermined maximum number of pages) and/or analyses of topics related to different parts of the course: 1) key privacy concepts; 2) privacy technologies; 3) privacy economics and policy. Homework assignments will be graded based on the clarity and validity of your arguments and analysis. You can use the **Objectives Rubric** attached at the end of this syllabus as a guideline to these criteria.

Some tips about how to complete a good assignment:

1. Think about the problem presented and described in the assignment.
2. Collect information about it.
3. Think about the problem again after you have studied the information collected.

4. Give a coherent structure to your ideas. Keep in mind clarity, concision, intellectual rigor, novelty, critical thinking, and analytical discussion.
5. Write down your ideas, remembering to cite collected information in a **properly formatted bibliography** (e.g., see http://www.chicagomanualofstyle.org/tools_citationguide.html). (Please note: while Wikipedia is a wonderful source of information it is NOT considered a *primary* source and should therefore not be cited in your work.)
6. Edit and proofread your text.

Please check carefully the schedule of HWs in the latter part of this document, and avoid scheduling meetings (e.g., job interviews) that conflict with your ability to submit the HW in time. While I understand that many of you may be interviewing soon for jobs, allowing students to submit HWs at a later time than the rest of their classmates would create unfair advantages. Hence, postponing HWs will not in general be allowed (and in the very exceptional circumstances in which it may be, the HW grading would be subject to heavy penalty).

3. The final project is a research project in which you try to address and answer an interesting (and hopefully novel) question related to or inspired by one of the topics covered in class.² The final project must be a group effort (from a minimum of 2 to a maximum of 4 students in each team). It can be “analysis”-oriented, or “coding”-oriented, depending on your skills and interests. For instance, it could be:

- An empirical study based on primary or secondary data (for instance, data that you mined/scraped/otherwise collected from the Internet).
- An experiment (for instance, a behavioral experiment).
- An analysis/critique of an existing privacy technology/framework/scenario.
- An analysis of the privacy implications/trade-offs an existing tool/app.
- An analysis of a new privacy attack vector.
- A novel privacy protecting (or invasive) app or tool of your design.
- And so forth.

On the other hand, as a general rule try and avoid projects of these types:

- Surveys of attitudes (i.e. projects that simply consist in asking people their opinions about some privacy sensitive issues).
- Literature reviews (i.e. mere reviews of what other people have studied/found out regarding certain privacy issues).

You can see the final project as a sort of small-scale research effort. To complete a good final project, you must come up with an interesting research question and a methodology, and produce some quantitative and/or qualitative results based on that methodology. The research question should at least strive to be novel and interesting, and the methodology should be sound and appropriate to investigating

² You will find out that the homework assignments and the final project build on one another: the HWs are designed to make you think about particular angles of the privacy debate which will turn useful in the completion of your project—and, hopefully, in other courses as well. In fact, you are encouraged to think about a project which builds upon one of the homework assignments.

that question. Therefore, literature reviews or reviews of a certain field are **discouraged** unless they are extensive, exhaustive, and novel.

The project, ideally, should be an original idea that you research (with proper citations of existing literature) within the constraints of a mini (that is, 7 weeks) course. It is suggested that you focus on a specifically defined topic rather than a broad field of enquiry (e.g.: “Re-identification of pseudonymous Twitter profiles using clickstream data” is much better than “Privacy in social networks”).

Your final project document should include the discussion and motivation for the research question you chose to investigate, a review of the relevant literature, a description of your methodology and/or hypotheses, and the analysis of your results. On average, a good final project document would be around 20 pages long and should be properly formatted according to academic guidelines (e.g., proper references and citations, structure and section titles, etc.).

Around mid-way through the course, you will be asked to select one specific idea and provide a proposal for your final project. The proposal should include your proposed topic, the research question you plan to investigate, your suggested methodology, your hypothesis and expected results, and an initial bibliography. The instructor and the TAs will give you feedback and suggestions based on your proposal. A list of possible project topics will be discussed in class, but you are strongly encouraged to propose your own topics to the instructor early enough in the course so to receive appropriate feedback – try to be creative!

The final project will be graded based on **creativity, novelty, and originality**, as well as on **how well you conduct research, your calculations and/or evaluations, the completeness of your work, your ability to draw conclusion, the clarity and presentation of your work, as well as its conciseness and organization**. Again, please use the Objectives Rubric attached at the end of this syllabus as a guideline to these criteria. In addition, students will be expected to present their results to the rest of the class during the last days of the course.

Note: final projects can be (and in the past often have been) later turned into independent studies, theses, or published articles. More generally, the skills needed to complete a proper project (and further emphasized below—for instance: the proper way to cite sources) will turn out to be useful to you in the completion of other Heinz/CMU courses, your final projects or Theses, and—hopefully—your future careers.

IRB approval for projects

If your research/project/case study involves experiments or surveys or other activity involving human subjects, or personal information about human subjects, you will need to understand concepts such as informed consent, human subjects’ protections, and IRB (Institutional Review Board) regulations before running your study. Most likely, you will also need approval for your study. Most typically, your study will follow under the “exempt” category, which requires you filling out a form with CMU IRB. Please check CMU IRB site: <https://www.cmu.edu/research-compliance/human-subjects-research/>.

Some tips about the final project

- You should **start thinking about the final project very soon** – this course only lasts seven weeks!
- I have uploaded to Canvas some projects from previous years that you can use as examples of interesting topics and analyses.
- You may find ideas for possible project topics by “glancing ahead” at the course’s readings throughout the various weeks.
- You may also find ideas by checking websites such as EPIC (consumer rights advocacy group: <http://www.epic.org/>) or IAPP (Privacy consortium: <http://www.privacyassociation.org/>); or by perusing additional academic papers available at <http://infoecon.net/workshop/bibliography.php>.
- Style guides are suggested rules for document formatting and reference notation. The use of a style guide is not required, but highly recommended if you intend to one day publish your work. (Even if you don’t, it makes formatting much easier). An APA style guide, quite common in social science, is described at <https://apastyle.apa.org/>.
- In order to offer some guidance for this research process, I have also uploaded **to Canvas a very rough draft document titled “How to do research” and another document about making a literature review.** These documents include a list of tips/hints/suggestions aimed at students who start a Master or a Dissertation research, but should also be useful for smaller projects like your privacy final project. For instance, the document offers information (as well as links to other resources) on how to
 - find and select an attractive research topic,
 - find information and look for references on that topic,
 - avoid plagiarism,
 - create a bibliography,
 - write a literature review,
 - organize your paper, and
 - write a research document.
- Here is **a rough timeline to guide you through the preparation of the final project:** after the first few weeks of the course, start brainstorming about the possible project topics. Think about which topics you are interested in; which topics are still worth studying (i.e., they have not yet been “over-harvested”); which topics are actually doable in the short time-span of a Mini; which topics may be of value to your future work/study at Heinz and to your future career. Send an email to me (the instructor) and the TAs to discuss your ideas and get feedback.

A NOTE ABOUT CITATIONS IN YOUR HWS AND PROJECTS

It is important to cite (in fact, *properly* cite) all sources in your HWs and final projects. Citations and references help you build on existing work (so you don't have to repeat what is already known), demonstrate your awareness of current literature, and support your arguments. They also help avoid you accidentally copying other people's work. As one previous TA of this class—Dr. Romanosky—wrote in his syllabus for a related class, proper citations and quotations allow you to avoid passing off someone else's work as your own: “[w]hether done accidentally or maliciously this is considered plagiarism and the academic equivalent of a crime. Proper use of quotations and citations will easily avoid this issue.”

ABOUT READINGS AND CLASS DISCUSSIONS

Classes will be a **combination of lectures and discussions**:

- Lectures: I will cover and extend the material and the topics covered in the readings.
- Discussions: we will also discuss together those topics – hence it is important that you know the readings before class.

Some classes have multiple readings; others, just one or two readings. Note that over the course of this Mini we will **not necessarily discuss each and every reading in class, neither will we cover in stringent detail each single reading**. Imagine readings and lectures as partially overlapping Venn sets: the classes touch upon the readings (but also touch upon materials *not* covered in the readings); and the readings offer you a means to learn more, and extend your understanding of, the same topics that we will cover in the slides.

TAKING CARE OF YOURSELF

Do take care of yourself – especially under these admittedly crazy times. Do your best to maintain a healthy lifestyle - eating well, exercising, getting enough sleep, and taking some time to relax. This will help you achieve your goals and cope with stress. Courses at CMU can be intense. If you are stressed out, please know that you are not alone, and that there are many helpful resources available on campus - an important part of the college experience is learning how to ask for help if it is needed. Asking for support sooner rather than later is often helpful. If you, or anyone you know, experiences academic stress, difficult life events, or feelings like anxiety or depression, please seek support: consider reaching out to a friend, faculty or family member you trust. Also, Counseling and Psychological Services (CaPS) is there to help: you can call 412-268-2922 or visit their website at <http://www.cmu.edu/counseling/>.

COURSE SCHEDULE AND TOPICS

Please note: In the schedule below, the readings are indicated with their full citation. On Canvas, you will find the respective files under the first author's last name and publication year. And did I mention that readings listed below should be completed **prior** to the class for which they are listed? That way we can discuss the related topics together during the lecture.

Also, please note that the first two weeks of readings are long, coalescing the discussion of the first week and providing background for future weeks. Please plan accordingly.

I hope you will enjoy and learn from this course!

Lecture 1 (Monday, August 31)

Topic: Introduction

Today's Reading: None

Lecture 2 (Wednesday, September 2)

Topics: Privacy, Security, and Anonymity

Today's Readings:

- Warren, S. D., & Brandeis, L. D. (1890). The right to privacy. Harvard law review, 193-220.
- Posner, R. A. (1978). Economic theory of privacy. Regulation, 2, 19.
- Solove, D. J. (2005). A taxonomy of privacy. U. Pa. L. Rev., 154, 477.

Monday, September 7

No classes – Labor day!

Lecture 3 (Wednesday, September 9)

Topic: Psychology and Behavior

Today's Readings:

- Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. Science, 347(6221), 509-514.
- Altman, I. (1977). Privacy regulation: Culturally universal or culturally specific?. Journal of social issues, 33(3), 66-84.
- Moore, A. D. (2003). Privacy: its meaning and value. American Philosophical Quarterly, 40(3), 215-227.
- Solove, D. J. (2020). The myth of the privacy paradox. Forthcoming.

Homework 1 uploaded to Canvas

Start proposing ideas for your projects to Alessandro and the TAs

Lecture 4 (Monday, September 14)

Topic: Privacy Intrusive Technologies - I

Today's Readings:

- Cranor, L. F., Sleeper, M., & Ur, B. (2013), Tracking and Surveillance. Mimeo.
- Froomkin, A. M. (1999). The death of privacy. Stan. L. Rev., 52, 1461. (**Only section I**)
- United States. Executive Office of the President, & Podesta, J. (2014). Big data: Seizing opportunities, preserving values. White House, Executive Office of the President.

Lecture 5 (Wednesday, September 16)

Topic: Privacy Intrusive Technologies - II

Today's Readings:

- Calo, R. (2011). The boundaries of privacy harm. Ind. LJ, 86, 1131.
- Narayanan, A., & Shmatikov, V. (2010). Myths and fallacies of "personally identifiable information". Communications of the ACM, 53(6), 24-26.
- Solove, D. J. (2007). I've got nothing to hide and other misunderstandings of privacy. San Diego L. Rev., 44, 745.

Lecture 6 (Monday, September 21)

Topic: Privacy Technologies - I

Today's Readings:

- Goldberg, I. (2007). Privacy enhancing technologies for the Internet III: Ten years later. Digital Privacy: Theory, Technology, and Practices, 3-18.
- Phillips, D. J. (1998). Cryptography, secrets, and structuring of trust. In: Technology and Privacy: The New Landscape. The MIT Press, Cambridge, MA, 243-276.
- Shen, Y., & Pearson, S. (2011). Privacy enhancing technologies: A review. HP Laboratories, 2739, 1-30.

Homework 1 due

Lecture 7 (Wednesday, September 23)

Topic: Privacy Technologies – II

Today's Readings:

- Brin, D. (1996). Excerpt from The Transparent Society. Wired magazine.
- Ohm, P. (2009). Broken promises of privacy: Responding to the surprising failure of anonymization. UCLA I. Rev., 57, 1701.

Homework 2 uploaded to Canvas

Finalize ideas for your projects with Alessandro and the TAs

Lecture 8 (Monday, September 28)

Topic: Privacy and the Legal System: US

Today's Readings:

- Bamberger, K. A., & Mulligan, D. K. (2010). Privacy on the Books and on the Ground. *Stan. L. Rev.*, 63, 247.
- Solove, D. J. (2016). A brief history of information privacy law. *Proskauer on privacy*, PLI.

Lecture 9 (Wednesday, September 30)

Topic: Comparative Legal Approaches to Privacy and Self-Regulation

Today's Readings:

- EU Court of Human Rights, Fact Sheet on the 'Right to be Forgotten' ruling. (C-131/12).
- Samuelson, P. (2000). Privacy as intellectual property?. *Stanford Law Review*, 1125-1173.

Lecture 10 (Monday, October 5)

Topic: Privacy and Economics I

Today's Readings:

- Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442-92.
- Einav, L., & Levin, J. (2014). Economics in the age of big data. *Science*, 346(6210).

Homework 2 due

Lecture 11 (Wednesday, October 7)

Topic: Privacy and Economics II

Today's Readings:

- Arrieta-Ibarra, I., Goff, L., Jiménez-Hernández, D., Lanier, J., & Weyl, E. G. (2018, May). Should We Treat Data as Labor? Moving beyond "Free". In *AEA Papers and Proceedings* (Vol. 108, pp. 38-42).
- Lenard, T. M., & Rubin, P. H. (2015). Big Data, Privacy and the Familiar Solutions. *JL Econ. & Pol.*, 11, 1.
- Rubin, P. H., & Lenard, T. M. (2001). Privacy and the commercial use of personal information.

Homework 3 uploaded to Canvas

Send draft of ongoing project to Alessandro and the TAs

Lecture 12 (Monday, October 12)

Topic: Privacy and Online Social Networks

Today's Readings:

- Marwick, A. E., & Boyd, D. (2014). Networked privacy: How teenagers negotiate context in social media. *New media & society*, 16(7), 1051-1067.
- Stutzman, F. D., Gross, R., & Acquisti, A. (2013). Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of privacy and confidentiality*, 4(2), 2.
- Wilson, R. E., Gosling, S. D., & Graham, L. T. (2012). A review of Facebook research in the social sciences. *Perspectives on psychological science*, 7(3), 203-220.

Lecture 13 (Wednesday, October 14)

Topic: Corporate Privacy and CPOs

Today's Readings:

- Bamberger, K. A., & Mulligan, D. K. (2011). New governance, chief privacy officers, and the corporate management of information privacy in the United States: An initial inquiry. *Law & Policy*, 33(4), 477-508.
- Spiekermann, S., & Cranor, L. F. (2008). Engineering privacy. *IEEE Transactions on software engineering*, 35(1), 67-82.

Lecture 14 (Monday, October 19) (This will be an interactive class, with students assigned to groups taking different sides of a policy debate)

Topic: Policy

Today's Readings:

- Pozen, D. E. (2016). Privacy-Privacy Tradeoffs. *The University of Chicago Law Review*, 221-247.

Homework 3 due

Week of October 19: Final project presentations (date TBA)

Friday, October 23: Final Projects due at noon

Objectives Rubric

Performance Element	Sophisticated	Competent	Not Competent
Conduct Research	Uses multiple sources, successfully locate appropriate and important information. When information is unavailable, explain the process done to try and locate the information.	Use multiple sources, successfully locate a few pieces of appropriate information but overlook at least one key piece of information. If not available, note that the necessary information is unavailable.	Use a source to locate some information successfully but overlook more than one key piece of information.
Calculations and/or Evaluation	Make all calculations and/or qualitative evaluations accurately. In-depth interpretation of results individually as well as an integrated whole. Makes logical assumptions but only when necessary and which are clearly explained.	Makes all calculations and/or qualitative evaluations accurately. Interpret results individually. Makes assumptions when necessary which are clearly explained.	Does not make all calculations and/or qualitative evaluations accurately. Limited interpretation of results OR make errors in calculations. Interpret results in isolation (individually). Makes illogical assumptions when not necessary or does not answer question because data is not readily available.
Completeness	Covers all questions. Proper discrimination used in including information on important related issued. No superfluous information included. Includes all proper attribution. When information is not available, a report of research done is provided. Reference to previous answers when necessary in answering subsequent questions with little or no repetition.	Covers all questions making it clear what issue is being addressed. Includes attribution. Explains when information is not available to answer question. Refers to parts of previous answers when necessary in answering subsequent questions.	Covers most questions but may answer multiple questions together without distinguishing which question is being addressed. Includes unnecessary information. Does not include reference to research.
Draw Conclusion	Uses information to draw logical conclusion. Supports conclusion with well thought out calculations and facts. Distinguish what the important facts are. Relate Analysis to topics discussed in class or in current events.	Uses information to draw logical conclusion but the conclusion is not clearly supported by data or some data is not accounted for by the conclusion.	Draws no conclusion or draws a conclusion that is not logical. Does not support conclusion with research and prior argument.
Clarity and Presentation of Written Expression	Expresses ideas clearly in own words, being careful not to duplicate presentation of the references; direct quotes held to a minimum. No spelling errors. Grammatically correct. Written report has professional appearance using graphics when appropriate which are referred to and explained in write up.	Expresses in own words, not duplicating language of the references; a few direct quotes but not excessive reliance on the language of research. Very few spelling or Grammatical errors but ideas are still clearly understandable. Written report has neat professional appearance.	Relies heavily on the words of references to present the argument. Many spelling & grammatical errors which interfere with readers' ability to understand. Written report does not appear professional. Includes no graphs, charts etc when appropriate or includes graphs and charts that are not referred to or are inappropriate.
Conciseness and Organization	Contains clearly developed ideas in a logical sequence. Completely answers all questions without unnecessary information. Refers to pertinent information already presented with little or no repetition of information.	Contains clearly developed ideas in a logical sequence. Completely answers all questions but may contain unnecessary information. Does not refers to pertinent information presented in earlier questions or part of analysis.	Is disorganized and difficult to follow. May not be complete. Repeats information unnecessarily.

[This page intentionally left blank]