

The Open INcentive Kit (OINK): Standardizing the Generation, Comparison, and Deployment of Incentive Systems

Noah Klugman
University of California, Berkeley
Berkeley, California
nklugman@berkeley.edu

Santiago Correa
University of Massachusetts, Amherst
Amherst, Massachusetts
scorreacardo@umass.edu

Pat Pannuto
University of California, Berkeley
Berkeley, California
ppannuto@berkeley.edu

Matthew Podolsky
University of California, Berkeley
Berkeley, California
podolsky@berkeley.edu

Jay Taneja
University of Massachusetts, Amherst
Amherst, Massachusetts
jtaneja@umass.edu

Prabal Dutta
University of California, Berkeley
Berkeley, California
prabal@berkeley.edu

ABSTRACT

Incentives are a key facet of human studies research, yet the state-of-the-art often designs and implements incentive systems in an ad-hoc, on-demand manner. We introduce the first vocabulary for formally describing incentive systems and develop a software infrastructure that enables UI-based graphical generation of complex, auditable, reliable, and reproducible incentive systems. We call this infrastructure the Open INcentive Kit (OINK). A review of recent literature from several communities finds that of the one hundred and twenty-one publications that incorporate incentives, only thirty-one describe their incentive system in detail, and all of these could be implemented using OINK. We evaluate OINK in practice by using it for an active energy monitoring deployment in Ghana and find that OINK successfully facilitates thousands of individual incentive payments. Finally, we describe our efforts to generalize OINK for different research communities, specifically focusing on architectural decisions around extensibility to support unanticipated use cases. OINK is free and open-source software.

CCS CONCEPTS

• **Human-centered computing** *Field studies*; • **Information systems** *Incentive schemes*.

KEYWORDS

Incentive System, Human Subject, Open Source

ACM Reference Format:

Noah Klugman, Santiago Correa, Pat Pannuto, Matthew Podolsky, Jay Taneja, and Prabal Dutta. 2019. The Open INcentive Kit (OINK): Standardizing the Generation, Comparison, and Deployment of Incentive Systems. In *Tenth International Conference on Information and Communication Technologies and Development (ICTD '19)*, January 4–7, 2019, Ahmedabad, India. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3287098.3287101>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ICTD '19, January 4–7, 2019, Ahmedabad, India
© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-6122-4/19/01.
<https://doi.org/10.1145/3287098.3287101>

1 INTRODUCTION

Experiments involving human subjects are common in a variety of fields, including psychology, economics, and computer science, among others. These experiments often encourage user participation by providing incentives to participants [14, 30, 50, 60]. It is known that incentive system design choices can have significant effects on the results of a study [5, 14, 22, 52], yet incentive systems receive relatively little discussion in the literature and, even when they do, there is no standard vocabulary to describe them [25, 52].

While information technology (IT) has revolutionized many aspects of the scientific enterprise—innovating classical measurement techniques such as surveys [24] and improving data analytics [55]—incentive systems have received little attention. Each experiment that uses incentives requires the researcher to develop methods to track when incentives are due, transfer the incentive, and provide accounting. In some cases, the cost of implementing a study-specific incentive system might be small—it may be easy to pay small populations for their participation in a survey—but implementation costs increase with the complexity of incentive triggers and with an increase in the size of the participant population.

To facilitate the description and comparison of incentive systems in the literature, and to reduce their experimental overhead while increasing their quality, we introduce the **Open INcentive Kit (OINK)**, the first work to propose a standard vocabulary for describing a generic incentive system and to support and automate the design, deployment, and management of end-to-end incentive systems.

OINK stems from the insight that very different studies share similar incentive system patterns. For example, a study that incentivizes a community health worker to visit a village (as reported from a mobile app tracking GPS) has a similar pattern to a study that incentivizes a driver to observe the speed limit (also as reported from a mobile app tracking GPS). OINK identifies and implements a core infrastructure that is common to many incentive systems, and offers a simple interface for implementing and sharing experiment-specific modules with other OINK users to be repurposed for their studies.

OINK supports high-quality, modular incentive systems for a range of simple to complex applications with minimal software programming. The system enables turn-key experiment planning and monitoring, including detailed transaction logging, visualizations of error states, and participation trends.

While incentive systems are used in a broad array of research settings, we initially introduce OINK as a tool for supporting research in development contexts, which have a history of innovative IT applications [28]. In particular, these settings challenge incentive systems because they have fewer sources of data or sensing devices and researchers may be working remotely with initially limited understanding of the context. These settings particularly benefit from the capabilities of OINK, including adapting to mobile money systems and simulating or modifying incentives to achieve a balance between effectiveness and manageability of incentives.

1.1 OINK Requirements

1.1.1 Reliable and Auditable. There is an ethical responsibility to provide all incentives promised and to ensure research funding is properly spent [11, 22]. An incentive system must always transfer incentives when expected. Further, it must provide bookkeeping to prove incentives are delivered as expected.

1.1.2 Broadly Accessible. The abstractions in this architecture must minimize the amount of work to generate incentive systems that fit the needs of a study. Anyone with a base level of technical literacy must be able to use the system without writing large amounts of code and maintaining a server.

1.1.3 Extendable. The architecture must allow for functionality to be developed by the community to fully suit a broad set of applications. Further, to ensure wide applicability, these extensions must be shareable and iterable between users.

1.1.4 Scalable. Incentive systems must be easy to scale both in population size and in complexity of incentive triggers. Further, the computing costs to run the system should be low regardless of scale.

1.1.5 Reproducible. When either the original designer or a third party replicates a study, the original incentive systems must be trivially recreated, shared, or directly reused.

1.1.6 Secure and Accessible. Incentive systems must prevent the unauthorized transfer of incentives stemming from either fraud or misconfiguration of the system. Additionally, incentive system administration must not require low-level server administration.

2 BACKGROUND

We consider work that explores the impact and regulation of incentive systems as well as work that demonstrates examples of IT being applied to human subjects research. OINK is the first work to explore principled generation and deployment of incentive systems.

2.1 Why Do Incentive Systems Matter?

Incentive system design is researched in many disciplines. The type of incentive chosen, the stimulus for incentive, and the method of incentive delivery have all been shown to have effects on presentation [14], selection bias [30], validity [9], response bias [50], as well as influencing other aspects of experimental design [5]. Thus, knowing more about an incentive system should be viewed as important in understanding a result [30]. OINK aims to make it easier for a community to include discussions of incentive systems design and implications by providing a common vocabulary for the literature.

OINK presents an opportunity to make incentive system design a first-order requirement of research projects. Auction-based mechanisms have shown in simulation that incentive design, including the amount, schedule, and transfer method, can significantly impact participant recruitment and retention, for both short-term and long-term participants [12, 27]. OINK enables principled characterization and iteration of these parameters.

2.2 Incentive Systems in the Literature

Stovel et al. examine 251 articles in five medical education journals to extract the incentives used [52]. They find 8% of papers directly describe the incentive system employed. This low percentage is presented as problematic due to both missed information transfer between researchers about incentive system design for future research and because it could hide underlying motivations of participants. We include the papers identified by Stovel et al. in our analysis in Section 4.1. We found no other work surveying the literature from other communities for descriptions of incentive systems used.

2.3 Incentive System Requirements and Policy

In the United States there are federal guidelines provided for human subject incentive systems [37], and universities adopt their own additional guidelines in their respective Institutional Review Boards (IRB) [11, 39]. Because there is no standard best practice, OINK does not guide users towards a specific ethical incentive system design; instead it leans on the standard practice of IRB submission, comments, and revision. A principled system like OINK opens up the possibility for the export of standard-language descriptions of any incentive systems constructed with it to be submitted to an IRB, potentially removing hurdles towards regulatory approval.

Similarly, governments enforce many different regulations regarding the electronic transfer of monetary funds [6, 8, 36, 38, 49]. OINK assumes that any digital money or tanagable good payment API invoked will bear any regulatory responsibilities.

2.4 Bringing IT to the Scientific Method

Information technology is already applied to human subjects experiments. A somewhat similar service to OINK is Amazon Mechanical Turk (MTurk), which provides a structured way to pay “Workers” for completing “Human Intelligence Tasks” (HIT) [1]. Unfortunately, MTurk does not allow for easy enrollment of participants and therefore could not easily be repurposed for a study that requires the enrollment of a specific population. Further, MTurk provides only the single stimuli (completing a HIT) which does not cover the full set of possible incentive systems. For example, if a study has a non-monetary incentive like providing a cookstove, MTurk does not provide a mechanism to facilitate this transfer.

IT is improving the generation and deployment of data collection systems, whether they are based on surveys [24], constrained geospatially [41], or incorporate sensors [7]. Open Data Kit (ODK), a suite of tools for data collection on mobile devices, has been adopted across and expanded by multiple academic and commercial communities, and is used everywhere from 42 countries to the International Space Station. [40]. OINK takes its name as an homage to ODK. We hope this work identifies a similarly powerful abstraction for incentive systems as ODK described for data collection.

3 DESIGN AND IMPLEMENTATION

Before designing OINK, we derive a generic incentive system vocabulary from a literature review (see Section 4.1). This vocabulary then informs a generic incentive system architecture, the implementation of which we call OINK. We explain how each of the words in our vocabulary is implemented in OINK and introduce the enabling technologies. Finally, we share wireframe designs of the OINK user interface to illustrate how we envision non-computer scientists interacting with the system.

3.1 Incentive System Vocabulary

We propose a vocabulary to describe the generic incentive system. To derive this vocabulary, we review 390 papers and abstracts to identify descriptions of incentive systems in any detail. We then identify the minimal vocabulary needed to describe each incentive system end-to-end. Finally, by taking the intersection of each of these descriptions, we can identify words and construct our vocabulary:

- **Stimulus:** This is the trigger event for a transfer of an incentive. It contains information about the incentive to be transferred on event and to whom. Example stimuli include the completion of a survey or the keeping of an app on a mobile phone for a period of time.
- **Payment:** Payment is the logic to transfer an incentive. Example payments include a call to the PayPal API or the manual transfer of a good to a participant.
- **Identity:** Identity is maintained for both participants and field officers (when applicable). Example usage include ensuring that incentives are transferred to the correct person using the correct method of transfer, tracking participant consent, and keeping track of the behaviors of the research staff either for fraud detection.
- **Fraud:** Fraud is the undesired transfer of an incentive. Fraud detection ensures the transfer of an incentive only after verification conditions has been reached. Example verification conditions include setting and enforcing a maximum number of hourly incentives per participant or stopping transfer of incentives to all participants located outside a pre-defined geographic area.
- **Alarm:** An alarm triggers if a condition within the incentive system is been met. Example conditions include payment error, system downtime, fraud, and budget exceeded.
- **Incentive System:** An incentive system is a system that keeps track of when a given amount of incentive should be sent (using a single stimulus and fraud detection), what incentive should be sent (using a payment), to whom an incentive should be sent, and whether an incentive was transferred (using an alarm on either completion or failure). It can only be triggered by a single stimulus and only causes a single transfer of incentive.
- **Experiment:** An experiment is a collection of one or more incentive systems that share common participants.

3.2 System Architecture

We introduce the OINK system architecture, which contains a runtime, high-level layers to support the administration of groups of incentive systems, and low-level layers that implement each word in the generic incentive systems described in Section 3.1. We then revisit the requirements introduced in Section 1.1 and describe the design decisions made to ensure each are satisfied. The OINK system hierarchy is shown in Figure 1 and described from the top down.

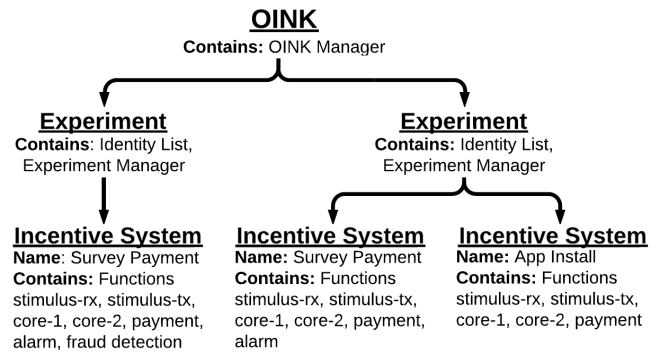


Figure 1: OINK system hierarchy. OINK sits at the highest level, followed by experiments and incentive systems. The OINK and experiment levels include managers, which provide administration, creation, and visualization at the respective level. Identity management is scoped to the experiment layer. Incentive systems are made up of multiple functions. Each function implements words from the generic incentive system vocabulary gathered from the literature. Notice that every incentive system is not the same, alarm and fraud detection functions are optional. All non-manager elements of the system run in a functions as a service (FaaS) cloud runtime.

3.2.1 Runtime. OINK is implemented using a functions as a service (FaaS) runtime. FaaS is a cloud-based architecture that gained commercial traction in 2014 [4]. It provides serverless computing, allowing computation to occur without configuring and hosting a long-running runtime environment. The primary programming abstraction for FaaS is a function, which is invoked, wakes up, computes, returns, and then is destroyed. Throughout the rest of the discussion, “function” refers to a function within a FaaS runtime.

The current OINK implementation uses the Google Cloud ecosystem, depending on Google Functions [16] for the FaaS architecture and Google Firestore for the datastore [15]. We chose to use the Google ecosystem for the OINK implementation because our first use case for OINK involved an Android app, which integrates well. Using Google tools for all components of OINK allows for a single authentication for all administration of the system. There is, however, no technical reason why OINK could not be entirely reimplemented using competing FaaS providers.

One limitation of FaaS is that it does not support long-running processes. Because of this, the OINK manager and the experiment manager—both of which are UIs and need to be accessible for as long they are being interacted with—are implemented as a webapp. The webapp authenticates with the FaaS runtime and datastore backend to load data and spawn new functions. If the webapp crashes, the FaaS runtime will continue hosting any active incentive systems.

3.2.2 OINK-Level: OINK Manager. This manages multiple experiments and is the entrance into the OINK UI. It allows environment variables to be set, including top-level authentication credentials required for administrating OINK, and provides an interface to create experiments. OINK can organize experiments with an arbitrarily large and complex number of incentive systems. The OINK manager is being implemented as a webapp and a preliminary wireframe design is shown in Figure 2a.

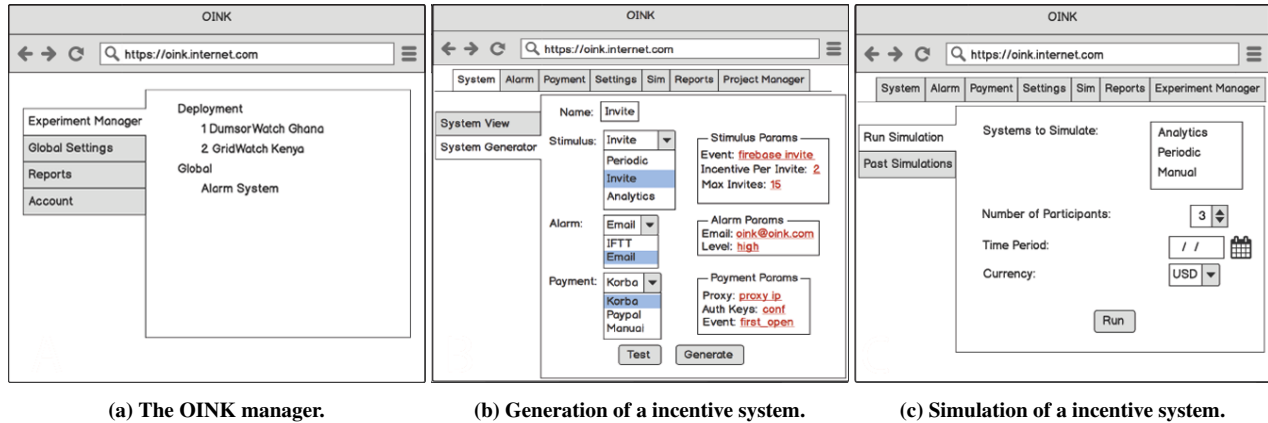


Figure 2: Three wireframes of the OINK system. (a) Shows the OINK manager with two available experiments (DumsorWatch Ghana and GridWatch Kenya). (b) Shows an experiment manager being used to generate an incentive system by selecting and parameterizing stimulus, payment, and alarm functions. This is generating the invite system (#2) from Table 2. (c) Shows the area in the experiment manager where the incentives consumed across multiple incentive systems can be calculated.

3.2.3 *Experiment-Level: Experiment Manager.* This contains multiple functionalities and is implemented as a page of a webapp. It is where a user generates incentive systems and is shown in Figure 2b and described in Section 3.2.5. It allows for high-level management of multiple incentive systems. For example, an experiment can have a budget, which if exceeded pauses all incentive systems in that experiment. It is where a user can simulate an experiment to see how much incentive will be consumed with different incentive system designs and is shown in Figure 2c. Finally, it is where visualizations of the performance of the incentive systems are located.

3.2.4 *Experiment-Level: Identity List.* Each experiment implements a single identity database that is used by all the incentive systems in the experiment. The identity database is default read, create, and update-only, which ensures consistency of participant identities across incentive systems. Each row in the identity table contains (at a minimum) a unique identifier, an array of valid payment methods, timestamps at participant creation and modification, and current participant status (enrolled, inactive, unenrolled).

3.2.5 *Incentive System-Level: Overview.* A single incentive system is shown in Figure 3 and each part is described in detail below. Every incentive system requires the OINK user to choose and parameterize exactly one stimulus function and at least one payment function. The core functions are included with every incentive system and require no parameterization. Along with these functions, multiple alarm and fraud detection functions can be added, each of which require selection and parameterization by the user.

3.2.6 *Incentive System-Level: Functions.* Every incentive system is implemented as a group of functions, most of which are also words defined within the generic incentive system vocabulary. A function can be understood to be a standalone logical block within the incentive system. We introduce all of the function types below and their interactions in Figure 3.

- **Core:** The core is the only area that does not contain words from the generic incentive system vocabulary. It is implemented by OINK and not modified by the user. It provides meta-functionality




Figure 3: Simplified Incentive System Diagram. This diagram shows an OINK incentive system. The *stimulus-tx* function is triggered to start the transfer of an incentive. The request moves to the *fraud detection* function which accesses historic transactions to determine validity. It then moves to the *core-1* function which logs the request. The *payment* function transfers the incentive. Then the *core-2* function stores the result and wakes up the *stimulus-rx* function, which performs stimulus-specific cleanup such as notifying the user. At any point, any function can trigger an *alarm* function.

for all incentive systems including logging and error handling. There are two functions and two databases that make up the Core: function *core-1*, function *core-2*, database *transaction-tx* and database *transaction-rx*. Function *core-1* wakes up on write to the *transaction-tx* table by a stimulus function. The *transaction-tx* database holds transaction information and includes all information needed to trigger a payment function. The *core-2* function wakes up and logs to the *transaction-rx* table when a callback from a payment function returns, and will invoke the *stimulus-rx* function for any stimulus-specific reporting.

- **Stimulus:** A stimulus is a trigger that results in incentive transfers. Each incentive system contains only a single stimulus. OINK provides an interface for a stimulus to connect to an incentive system. For example, OINK provides a way for someone who has gathered a survey response to automatically trigger an incentive. The parameters for this stimulus would be the person to pay, the person who conducted the survey, the amount of money to pay, the desired method of payment, and when to pay (immediately, after upload and fraud check, etc.). This stimulus then exists over the lifetime of the incentive system and allows for the field officer to trigger an incentive transfer by collecting and transferring this data to OINK.

There are two functions required to implement each stimulus: *stimulus-tx* and *stimulus-rx*. The *stimulus-tx* function contains the logic to trigger OINK, the amount of an incentive that this trigger should release, who the incentive should be transferred to (as an index to the experiment-level identity table), how the incentive should be transferred, and how the result of a stimulus should be displayed back to the user. The *stimulus-rx* function writes the result of a payment into the *transaction-rx* database and optionally does stimuli specific clean up. For example, a stimuli generated from an action in a mobile app action might want to display a message back to the user after a successful payment in that app. Functions implementing a small group of stimuli, including an HTTP REST API, a cron-based periodic system, and a Google analytics receiver [19], are provided in the OINK codebase to demonstrate standard ways of triggering the incentive system.

- **Payment** A *payment* function provides logic to transfer an incentive to a user. For example, a *payment* function could be a wrapper to the third party payment API. As there is a one-to-one mapping of a stimulus to a payment, there is only a single *payment* function per incentive system. The inputs to a *payment* function include the amount of the incentive, the type of incentive, the time a payment was attempted, the reason an incentive is being generated, and the identity of the receiver. The output from a *payment* function is a confirmation that a given incentive transfer occurred or an error message and is collected by the *core-2* function and written to the *transaction-rx* table. Additionally, the *payment* function can be parameterized as a pass-through, allowing for an experimenter to opt-out of transferring incentives, which can be useful for testing. A starting group of *payment* functions including the Korba airtime distribution API [26] and a manual payment over SMS are implemented in the OINK codebase.
- **Fraud Detection** The *fraud detection* function contains rules that can prevent a stimulus from triggering an incentive in order to prevent unattended transfer. For example, a *fraud detection* function can ensure that a participant receives no more than a maximum number of an incentive. If this is violated, the suspect transaction is paused until a researcher manually investigates. Fraud detection functions can access any of the databases in OINK to query historic state for any type of undesired patterns. Multiple types of *fraud detection* functions can be defined within a single incentive system.
- **Alarm:** Any part of an incentive system can trigger an alarm. The input to an *alarm* function is a message and the output can be an API, webhook, or an internally-provided function to generate and transmit a notification. *Alarm* functions are useful for error tracking and for periodic updates. Multiple *alarm* functions can be included per incentive system. A small group of *alarm* functions are implemented in the OINK codebase including email and SMS [56].

3.2.7 Revisiting Requirements. We revisit each system requirement in Section 1.1 and describe how these requirements are satisfied by the OINK design and implementation.

3.2.8 Requirement: Reliable and Auditable. We implement OINK on top of the Google FaaS runtime [16]. This bootstraps reliability off Google who, like each of three largest commercially-available FaaS runtime providers, employ dedicated staff to keep their FaaS runtime operational. Recognizing that the FaaS runtime creates a log each time a function is invoked and destroyed, we design OINK to

invoke a new function whenever data is moved through an incentive system. This ensures a record at the entry and exit point of each stimuli, fraud, payment, core, and alarm function.

3.2.9 Requirement: Broadly Accessible. The FaaS runtime by definition removes the need for the configuration and maintenance of a long-running server to host OINK, reducing the barrier of entry for non-technical users. Further, we designed OINK with input from non-computer science researchers and the first deployment of OINK supports an interdisciplinary project. We received early feedback on the usability of OINK based on wireframes of the user interface (shown in Figure 2) and will focus on usability as future work.

3.2.10 Requirement: Expandable. To support large numbers of different incentive system configurations, OINK is designed so that each word in the incentive system can be completely customized. This is done by coding a function for that word, which can be a technically challenging task. Thankfully, once a custom function is written, the FaaS runtime allows for it to be shared between users to be incorporated into their incentive system design. We hope that once popular functions have been implemented once, it will be unlikely that a user will have to implement an extension to OINK. We are motivated by the emergence of people sharing useful FaaS functions to support applications on other platforms [34].

3.2.11 Requirement: Scalable. Every commercially available FaaS runtime operates at scale and supports millions of function invocations out-of-the-box. This allows OINK to scale without user intervention. As more incentives are transferred, more functions will simply be called. Because every major FaaS platform allows at least one million free invocations monthly [2, 16, 31], OINK will experience minimal operating costs for many experiments.

3.2.12 Requirement: Reproducible. OINK can be reproduced without requiring any low-level server configuration. This is possible both because the FaaS runtime will be the same for all users, and because all major FaaS providers provide databases and virtual servers that can be launched pre-configured to host the OINK databases, OINK manager and experiment manager webapps. Further, each part of OINK is independently reproducible because the same function will run on any FaaS runtime (requiring a lightweight adaptation layer to the specific FaaS runtime environment).

3.2.13 Requirement: Secure and Accessible. OINK bootstraps security and access control off the authentication system of both the FaaS service and the payment API. Each of the three FaaS services evaluated provide abstractions to protect and provide levels of access [2, 16, 31]. By keeping all functions and logs within this system, data generated and the code being run are secured by the FaaS service provider. Some stimuli themselves can bootstrap security with these same protections; for example, an app developed with Firebase Cloud Messaging can use Google APIs to authenticate with an OINK system using Google Functions as its FaaS backend out-of-the-box [17]. For monetary or other digitally transferred incentives such as airtime, OINK bootstraps security off methods implemented by the external APIs performing the transfer.

Type	Survey			Deployment			Both
	No Mention	Some Mention	Discussed	No Mention	Some Mention	Discussed	No Human
Papers	15	11	22	19	18	5	27
Abstr	8	1	2	9	9	2	36
Total	23	12	24	28	27	7	63

Table 1: Discussion of incentive systems found in 390 publications: 150 ACM DEV papers and shorts, 215 medical education research papers, 3 ICTD Papers, and 22 CHI Papers. Note the small number of papers that discuss incentive systems in detail. For each paper that discusses an incentive system, we find that the OINK incentive system vocabulary fully describes the incentive system.

4 EVALUATION

We evaluate incentive systems in the wild by performing a literature review of 175 papers across three development-related ACM communities. We also consider work containing descriptions of incentive systems discovered from an examination of 215 publications from medical journals [52]. We evaluate the performance of OINK under an ongoing deployment managed by the authors where participants are incentivized for nine different types of activities. OINK is being used to implement all incentive systems for this deployment and has issued over three thousand incentive transfers. Finally, we select the most complicated incentive system found in our literature review — taken from *A Feasibility Study of an In-the-Wild Experimental Public Access WiFi Network* which appeared in the fifth ACM DEV — and discuss how OINK could be configured to support this work [48].

4.1 Incentive Systems in the Wild

To uncover the taxonomy of incentive systems in the technology for development community, we perform a literature search.¹ We select papers using the Google Scholar query "source:Computing source:Development" and select the first 150 papers [21]. We review each paper and decide whether they had no incentive system (papers using existing datasets, wireless network protocols, etc), or if an incentive system was present. We review prior CHI and ICTD proceedings based on titles that seem likely to have human factors research. Finally, we review papers identified by Stovel et al. that contain incentive system descriptions from the medical community [52]. Because of the intuition that a survey might have a simpler incentive system that would be less likely to be reported in the literature, we then classify each paper's incentive system as either triggered by a survey or by a deployment. This line can be fuzzy at times and for these cases, we make a subjective decision and place the work in a single category. We further categorize works based on their discussions of the incentive system: no mention, some mention, or discussion. When a paper mentions giving someone technology or asking someone questions but does not go further, we consider this "no mention." When a paper mentions the work being done in a school or implies that participants volunteered, we classify as "some mention." When a paper mentions what people received or what triggered an incentive, we classify as "discussion." The results of this classification are summarized in Table 1.

¹The full citation of the papers selected and their classifications can be found at <https://www.github.com/lab11/OINK/papers/ictd18/citations.csv>.

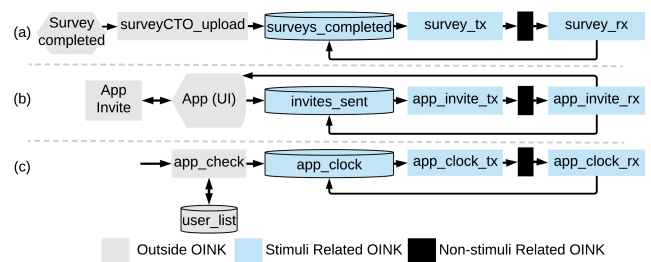


Figure 4: Three different stimuli from DumsorWatch. In (a) a completed survey is uploaded which writes to a collection (surveys_completed) and wakes up the survey_tx *stimuli-tx* function. In (b) a user sends a download invitation in an app, writing to a collection (invites_sent) and waking up the app_invite_tx *stimuli-tx* function. Finally, in (c) a cron job (app_check) wakes up on a timer, checks to see how long users from the identity table have been active, and then writes users active longer than a threshold to a collection (app_clock) waking up the app_clock_tx *stimuli-tx* function. The grey shapes show logic that occurs outside of OINK and demonstrate that different stimuli interface with OINK using a common interface, a single write to a collection. The blue shapes show the OINK stimulus logic, and the black box represents the rest of the incentive system. Mapping these to Table 2, (a) is #9: surveyCTO, (b) is #2: app_invite, and (c) is #3 or #5: keeping either the app (app_clock) or sensor (sen_clock) installed for a period of time.

From Table 1 we can make a couple of observations. Broadly, we find a majority of papers with incentive systems do not describe these systems either at all or beyond trivial details. Of all papers found involving human subjects, 32% contain at least a minimum amount of information about their incentive system and 26% speak in detail. Grouping this by communities, we see that out of the CHI and ICTD papers (selected for our review based on titles suggesting human studies), 84% had human subject studies. Of these, 28.5% mention their incentive system and 42.8% have a discussion of the incentive offered. We discover 59 papers (39%) from the set of ACM DEV papers (which were selected for our review regardless of title) that have some amount of human subject work 21% mentioned incentive systems and only 6% thoroughly described incentive systems. These are lower percentages than the nearest related work, which found 8% of work directly describing incentive systems [52].

4.2 Scenario 1: DumsorWatch Pilot

DumsorWatch is an experiment measuring the reliability of the distribution level of the AC power grid in Accra, Ghana. Participants consent and are surveyed in-person, a mobile app is installed on participants' everyday-use Android smartphones, and grid reliability sensors are installed in some households. Participants are incentivized at different levels for these different activities. All incentives for participation are provided as airtime top-ups.

4.2.1 Experiment Walkthrough. Table 2 provides an overview and explanation of all incentive systems within the experiment, and Figure 4 demonstrates some of the stimuli in detail. Overall, the DumsorWatch experiment demonstrates a real-world example of the breadth and complexity of the incentive systems enabled by OINK.

	1	2	3	4	5	6	7	8	9
name:	app_install	app_invite	app_behavior	sen_install	sen_behavior	sen_collect	debug	FO_manual	survey
stimulus-tx:	first_open	app_invite	app_clock	sen_query	sen_clock	sen_pick	debug	FO_manual	surveyCTO
fraud detection:	user_thresh	user_thresh	none	none	none	user_thresh	none	user_thresh, FO_thresh	user_thresh, FO_thresh
payment:	Korba	Korba	Korba	Korba	Korba	Korba	none	SMS	Korba
stimulus-rx:	first_open	app_invite	app_clock	sen_query	sen_clock	sen_pick	debug	FO_manual	surveyCTO
alarm:	email	email	email	email	email	email	email	email	email

Table 2: The DumsorWatch experiment incentive systems described by the OINK vocabulary. Each column is a single incentive system used in the DumsorWatch deployment. Each row is a function hosted with the FaaS architecture. The incentive systems send airtime when: (1) installing an app for the first time, (2) a user invites a friend to install the app, (3) the user keeps the app installed, (4) a sensor is installed in a household, (5) a sensor is kept in a household, (6) a sensor is collected, (9) a survey is completed. Additionally, (7) allows for debugging and (8) lets a field officer manually transfer airtime for unanticipated reasons. The `user_thresh` and `FO_thresh` fraud detection functions ensure no more than a maximum number transactions occur for a user or are generated by a field officer.

4.2.2 Incentive Walkthrough. We take a closer look at a single incentive system: `app_invite` (incentive system 4 in Table 2). This incentive system transfers credit to a user when they use the DumsorWatch app to invite an additional participant to download the app. For each function we discuss its implementation and break out its parameters, which are set by the researcher based on the specific requirements of the incentive system.

Stimulus-TX: Firebase Invites provide a turn-key mechanism for sending download invitations to contacts in a phone either over email or SMS [18]. When the user triggers an invite in the app, the DumsorWatch app writes to a Firestore collection using the Firebase Android library [15]. An `on_create` trigger on this table wakes up the `app_invite` *stimulus-tx* function. This function contains the logic of the identity of the person receiving the incentive, the method of transferring the incentive, the amount to pay per invite and the threshold number of invites allowed per user.

Fraud: The *fraud* function detects if a user has triggered an amount of invites greater than allowed. If this threshold is not hit, the fraud block transfers data to the core and marks the stimulus collection "enqueued". If the *fraud* function rejects a transaction, it marks the stimulus database as "fraud" and stops the transfer of data until manual intervention by the researcher. It then triggers an *alarm* function with information about the fraud.

Core-1: The *core-1* function receives a payment record and is responsible for logging, error handling, retries, and guaranteeing that the record will be updated upon success or failure of incentive delivery. It has no parameters and is in all incentive systems.

Payment: This incentive system uses the Korba service to deliver incentives. The Korba *payment* function is responsible for mapping the generic payment information from OINK (whom and how much) to the format expected by the Korba API [26]. In practice, the Korba API exposed a limitation of the FaaS architecture. For security, Korba requires that all payment requests come from a whitelisted IP, and FaaS systems do not support static IPs. As a workaround, we deploy a dedicated proxy server for OINK requests to Korba.

Core-2: The Korba API includes a callback URL to invoke the *core-2* function upon payment completion or failure. *Core-2* updates the transaction database with the API result and invokes a *stimulus-rx* function. It has no parameters and is in all incentive systems.

Stimulus-rx: This function is responsible for handling the completion (or failure) of a *payment* function. For this incentive system, the *stimulus-rx* is configured to retry up to five times, after which

point payment is paused and the *alarm* function is called to notify the research team. For the `app_invite` incentive system we want to notify the user that sent an invite that their payment has been disbursed, so the *stimulus-rx* function spawns a message back to the app using Firebase Cloud Messaging [17] to be displayed in a notification.

Alarm: The *alarm* function takes a message and sends an email to the project team. Eventually, we envision a more nuanced set of alarms, with varying priorities and mechanisms.

4.2.3 Other Incentive Systems in DumsorWatch. We briefly discuss parts of the other incentive systems in DumsorWatch, each of which is shown in Table 2. In (1) the *stimulus-tx* function `first_open` triggers when an app is first installed. The *fraud detection* function `user_thresh` only allows a user to receive a "first install" incentive once by looking at historic data on the user. The *payment* function `Korba` uses the Korba Payment API to send airtime topup [26]. On a successful payment response from the Korba API, the `first_open` *stimulus-rx* function sends a confirmation message to back to the app. The *alarm* function `email` sends an email if any part of this chain fails. (3) and (6) both use cron-jobs to pay incentives after users have kept either the app or the sensor installed for a period of time. (7) shows a configuration that allows for the ad-hoc manual transfer of an incentive in a principled manner, ensuring that this transfer is logged. (8) and (9) use two fraud-detection functions; one that checks the recipient and one that checks to behavior of the field officer (i.e. don't let a field officer report more than N surveys a day). Finally (8) demonstrates how a field officer can send an SMS to create a record in OINK after manually transferring an incentive.

4.2.4 Field Performance. OINK has transferred nearly three thousand incentives during three months of deployment. It has had experienced 100% uptime. Further, it has proven to be fault resilient in other ways, with the alarm functions letting the team know quickly when there are problems (a common problem turns out to be the reliability of the Korba API) by sending emails with a targeted error message and an index to trace the lower-level logs. Over the course of the deployment there have been multiple reasons to audit particular payments, each of which OINK has alerted the research team about and provided sufficient logs to address. The research group is interdisciplinary and includes an graduate student in economics, who has been able to interact with OINK for day-to-day management activities. An upcoming large-scale DumsorWatch experiment will continue to use OINK for incentives.




Figure 5: Incentive System from “A Feasibility Study of an in-the-wild Experimental Public Access WiFi network” [48]. One design using OINK to support the incentive system described in [48], which transfers vouchers to participants based on the amount of WiFi bandwidth they share.

4.3 Scenario 2: A Paper from the ACM DEV Literature

We infer the design of an incentive system from the description in “A Feasibility Study of an in-the-wild Experimental Public Access WiFi network” [48]. This paper explores constraints in providing free public WiFi network — created by recruiting participants to sharing their private broadband connections — in a low income neighborhood in a medium-sized British city. We choose this paper from our survey of incentive systems because it contains the most complicated incentive system we were able to find. Figure 5 shows our mapping of this incentive system to OINK. The authors describe their incentive system as follows:

For their time and the inconvenience, participants were offered compensation in the form of shopping vouchers. . . The size of compensation was scaled in line with the inconvenience suffered: sharers who made their broadband available to others. . . received higher compensation (100 pounds) than those citizens who simply used the. . . service (50 pounds).

4.3.1 Incentive Walkthrough. One possible incentive system design supporting this work imagines a web-form as a stimuli. A field officer enters the participant ID, their ID, if a participant shared their broadband, and how much they shared. When the web-form is submitted, it generates a HTTP POST, which OINK receives using a *stimulus-tx* function and passes to the *payment* function. The *payment* function then generates a SMS message to the field officer with the amount of voucher to transfer. To finish, the field officer enters the participant ID, their ID, and whether the incentive was transferred into the web-form. On submit, OINK receives another HTTP POST using a *stimulus-rx* function and marks the transaction complete. This full incentive system is shown in Figure 5.

4.3.2 What Could OINK Have Added? This paper does not talk about incentive management. We imagine a basic *fraud* function to ensure that no user is paid twice and an *alarm* function that sends an email if the transfer of incentives fails or if fraud is detected.

OINK allows for higher frequency payments than the authors used. For example, the experiment only paid people at the end. Perhaps multiple midpoint payments would have improved participation rates by generating periodic reinforcement of good behavior [5, 12]. Similarly, OINK allows for higher resolution stimuli. For example, the experiment also only had two voucher levels; more incentive levels would allow the exploration of whether an incentive level within reason exists that would lead to a non-linear increase of users willing to provide public access WiFi [12, 58].

5 DISCUSSION AND FUTURE WORK

5.1 Architecture Limitations

5.1.1 New Features Require Custom Code. Stimuli and payments require a fair amount of software development to first be incorporated into OINK, which will lead to technical hurdles for early adopters. For example, the app-invite described in Table 2 contains over a hundred lines of code written specifically for that stimulus. Similarly, the Korba payment function took considerable engineering effort learning and debugging with the Korba API. We imagine that as OINK matures, a set of templates will emerge to simplify this process. To bootstrap this, we provide example stimuli, payment, fraud, and alarm functions in the publicly available codebase that demonstrate our best practices. A partial list of these include a SurveyCTO stimuli [54], a Google Analytics stimuli [19], a SMS based stimuli, a manual stimuli, a manual payment, a Korba payment [26], and an email alarm. OINK is architected such that once this work has been done, these functions can be shared with a community to be reused across experiments with minimal reconfiguration.

5.1.2 Static IPs and Whitelisting. None of the three commercial FaaS services currently provide a static IP service, although it is a commonly requested feature [10]. This is a problem when using third party APIs that use whitelisting in their security systems. For our OINK implementation of the Korba API (which uses whitelists), we configure the *payment* function to call a whitelisted server rather than the Korba API directly. This server then forwards the request to the Korba API and the response back to OINK. Requiring an always on whitelisted proxy server violates a true serverless architecture.

5.1.3 Latency. All three FaaS providers have latency guarantees in milliseconds, which could be too slow for some applications.

5.1.4 Infinite Loops. None of the three FaaS providers have explicit protection against infinite loops created by functions that call themselves. Some protection is provided by a hard limit in the number of functions that can be run in a unit of time, but this can lead to unexpectedly high FaaS costs. While this is a danger, it can be protected against with careful software development and testing.

5.2 Incentive Systems As Their Own Intervention

5.2.1 Communicating Incentive System Design Patterns. The argument can be made that because incentive systems can impact the results of an experiment, each experiment that involves incentives should report the design of their system [52]. In our literature review we find that only 30% of papers include descriptions. Perhaps the barrier towards including more incentive system descriptions will be lowered by our definition of a standard vocabulary to describe incentive systems. Or, perhaps, OINK could export language-based descriptions of the incentive systems running in an experiment to be trivially included in the literature.

Similarly, Computer Science has been working towards adherence to Institutional Review Board (IRB) approval as a requirement for publication [13]. We aim to work with different IRB’s to establish an method of automatically exporting descriptions of experiments designed using OINK in a standardized language-based format that could be presented as part of an IRB protocol.

5.2.2 Experimental Design. Both by increasing the resolution of stimuli and by allowing tighter control of the frequency of incentive transfer, OINK could enable previously difficult to conduct experiments. For example, if a researcher was interested in incentivizing a user to check their blindspots more while driving, OINK could allow stimuli from an eye tracker. Or, a researcher could study how incentives influence behavior over the course of a day by using OINK to provide different amounts of credit for the same action over the course of the day. Programmers could generate code with relative ease implementing either of these examples, but OINK would allow people less comfortable coding to run these experiments.

5.2.3 Simulation. The principled implementation of an incentive system opens up an opportunity for tools to be developed that allow for the quick simulation and iteration of incentive systems. These tools could be helpful to scope different experimental designs when faced with budgetary constraints.

5.2.4 Incentive System Research. Many crowd-sourcing works explore incentive systems design to attract and hold crowds of specific sizes or expertise [59]. OINK can ease the implementation and evaluation of these novel incentive systems.

5.3 Other Uses for OINK

OINK can be extended to many different domains. For example, OINK could distribute bonuses at a company by using key performance indicators as a *stimuli* function, and payroll as a *payment* function. It could be used to track non-tangible incentives using a design where a public health worker uses OINK to email words of encouragement every time a patient arrives at a clinic on time by using the arrival time as a *stimuli* function, and an email service as a *payment* function. Or, it could be used for non-human scientific studies by dispensing a food when a button is pressed (a common incentive when studying animals [29, 45, 47]) with a button press as a *stimuli* function and actuator control as a *payment* function.

5.4 Future Work

5.4.1 Full Release. OINK is in an alpha state; the mature parts of the implementation are those required to support the DumsorWatch deployment. It is complete enough for experienced computer scientists to deploy, but not yet a turn-key solution for domain scientists.

5.4.2 Community. To minimize development effort, OINK users should have access to multiple pre-implemented functions that support the majority of incentive systems. Some functions are provided in the OINK code base, but many more will have to be implemented. We plan on hosting OINK tutorials to attract a community of early adopters to contribute to this effort.

5.4.3 Usability Testing. The OINK system has been designed based on input from by a small group of researchers. We will seek usability feedback from larger and more diverse groups of users.

5.4.4 Cross-community Literature Review. OINK aims to be a tool for many different scientific communities. It has been implemented to support the communities we are most familiar with. However, we expect branching out will reveal edge cases not previously encountered, potentially requiring revisions to the OINK architecture and our incentive system vocabulary.

5.4.5 Expansions. We enumerate a list of potentially useful extensions for OINK left for future development. Possible *stimuli* expansions include libraries that respond to events on mobile devices [3, 18, 33], web user interfaces triggered by user interaction [20], and events generated from data-streams collected by sensors [43]. Possible *payment* expansions include support for different mobile money APIs [6, 44, 51, 53], non-monetary non-tangible incentives such as sending a supportive email [35, 52], or non-monetary tangible incentives such as releasing food to an animal [45, 46]. Possible *alarm* expansions include services that generate emails [35], SMS messages [56], and graphical displays [23]. Possible *fraud* expansions include machine learning services to detect anomalies [32] or provide methods for additional authentication [57].

One notable external service, especially for alarms and stimuli, is IFTTT, which allows for the graphical construction of applets that incorporate multiple APIs [42]. For example, IFTTT can automate the configuration of a system that sends an HTTP POST packet when a phone enters a location. This HTTP POST could then be used as an OINK stimulus to trigger an incentive, allowing for construction of a complicated incentive system with writing minimal (if any) code. Similarly, IFTTT can respond to incoming web traffic, allowing OINK alarms to integrate easily with third-party APIs.

5.4.6 Implementation Using Other FaaS Implementations. There are three major FaaS providers who could support OINK [2, 16, 31]. A researcher using services from one cloud may want to run OINK in that same environment. Porting OINK to different FaaS environments remains future work.

6 CONCLUSIONS

Incentive systems are critical to human subjects research, often impacting the results of the experiments themselves. Their architectures, however, are not well discussed in the literature, and current incentive systems are complicated to construct, difficult to scale, and limited without integration with a large number of stimuli and payment APIs. We claim that these challenges impede human subjects research. To address this problem we present the Open INcentive Kit (OINK), a software system that provides out-of-the-box automation, autogeneration, scalability, security, and accountability for incentive management. We believe that OINK can add structure to the incentive design process, potentially leading to broad adoption by human subjects researchers. This enables incentive systems that can be shared and reused across a community, ultimately modernizing this critical component of human subjects research.

7 ACKNOWLEDGMENTS

This work was supported in part by the Development Impact Lab (USAID Cooperative Agreement AID-OAA-A-13-00002), part of the USAID Higher Education Solutions Network. We thank Joshua Adkins for his technical guidance, Meghan Clark, Zakir Durumeric, William Huang, and Lane Powell for their edits on this paper, Da Seglah and Nana Bonku Eshun for their technical support with Korba, Kwame Abrokwah, Susanna Berkouwer, and Catherine Wolfram for being early adopters of OINK, the anonymous reviewers for their insightful feedback and our anonymous shepherd for their guidance.

REFERENCES

- [1] Amazon. 2018. Amazon Mechanical Turk. <https://www.mturk.com/>. (2018). (Accessed on 07/22/2018).
- [2] Amazon. 2018. AWS Lambda – Serverless Compute – Amazon Web Services. <https://aws.amazon.com/lambda/>. (2018). (Accessed on 02/28/2018).
- [3] Amazon. 2018. AWS Mobile. <https://aws.amazon.com/mobile/>. (2018). (Accessed on 03/02/2018).
- [4] Amazon. 2018. Document History - AWS Lambda. <https://docs.aws.amazon.com/lambda/latest/dg/history.html>. (2018). (Accessed on 03/02/2018).
- [5] Colin F Camerer and Robin M Hogarth. 1999. The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of risk and uncertainty* (1999).
- [6] Indian Government Cashless India. 2018. Unified Payments Interface (UPI). <http://cashlessindia.gov.in/upi.html>. (2018). (Accessed on 03/02/2018).
- [7] Rohit Chaudhri, Waylon Brunette, Mayank Goel, Rita Sodr, Jaylen VanOrden, Michael Falcone, and Gaetano Borriello. 2012. Open Data Kit Sensors: Mobile Data Collection with Wired and Wireless Sensors. In *Proceedings of the 2nd ACM Symposium on Computing for Development*. ACM, 9.
- [8] PCI Security Standards Council. 2018. Official PCI Security Standards Council Site. <https://www.pcisecuritystandards.org/>. (2018). (Accessed on 03/02/2018).
- [9] Robin P Cubitt, Chris Starmer, and Robert Sugden. 1998. On the validity of the random lottery incentive system. *Experimental Economics* 1, 2 (1998), 115–131.
- [10] James Daniels. 2017. Google Cloud Functions static IP. <https://groups.google.com/forum/#!topic/firebase-talk/o9Br8GtYBN4>. (2017). (Accessed on 03/02/2018).
- [11] Committee for Protection of Human Subjects. 2017. Compensation of Research Subjects. <https://cphs.berkeley.edu/compensation.pdf>. (2017). (Accessed on 02/28/2018).
- [12] Lin Gao, Fen Hou, and Jianwei Huang. 2015. Providing long-term participation incentive in participatory sensing. In *Computer Communications (INFOCOM), 2015 IEEE Conference on*. IEEE, 2803–2811.
- [13] Simson L Garfinkel and Lorrie F Cranor. 2010. *Institutional review boards and your research: a proposal for improving the review procedures for research projects that involve human subjects and their associated identifiable private information*. Technical Report. Naval Postgraduate School.
- [14] Uri Gneezy, Stephan Meier, and Pedro Rey-Biel. 2011. When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives* 25, 4 (2011), 191–210.
- [15] Google. 2018. Cloud Firestore. <https://firebase.google.com/docs/firestore/>. (2018). (Accessed on 03/01/2018).
- [16] Google. 2018. Cloud Functions. <https://cloud.google.com/functions/>. (2018). (Accessed on 02/28/2018).
- [17] Google. 2018. Firebase Cloud Messaging. <https://firebase.google.com/docs/cloud-messaging/>. (2018). (Accessed on 03/01/2018).
- [18] Google. 2018. Firebase Invites. <https://firebase.google.com/docs/invites/>. (2018). (Accessed on 03/01/2018).
- [19] Google. 2018. Google Analytics Solutions - Marketing Analytics & Measurement. https://www.google.com/analytics/#?modal_active=none. (2018). (Accessed on 03/02/2018).
- [20] Google. 2018. Google Forms App Script. <https://developers.google.com/apps-script/reference/forms/>. (2018). (Accessed on 03/02/2018).
- [21] Google. 2018. Google Scholar. <https://scholar.google.com/>. (2018). (Accessed on 03/02/2018).
- [22] Christine Grady. 2005. Payment of clinical research subjects. *The Journal of Clinical Investigation* 115, 7 (2005), 1681–1687.
- [23] Grafana. 2018. Grafana - The open platform for analytics and monitoring. <https://grafana.com/>. (2018). (Accessed on 03/02/2018).
- [24] Carl Hartung, Adam Lerer, Yaw Anokwa, Clint Tseng, Waylon Brunette, and Gaetano Borriello. 2010. Open Data Kit: Tools to Build Information Services for Developing Regions. In *Proceedings of the 4th ACM/IEEE international conference on information and communication technologies and development*. ACM, 18.
- [25] Matthew Jadud. 2016. IRB Reviews Required. (2016). <https://doi.org/10.1145/2993223.2993229>
- [26] Korba. 2018. Korba - Our Lives Simplified. <http://korbaweb.com/library>. (2018). (Accessed on 03/01/2018).
- [27] Iordanis Koutsopoulos. 2013. Optimal incentive-driven design of participatory sensing systems. In *Infocom, 2013 proceedings ieee*. IEEE, 1402–1410.
- [28] Stanley Lippert. 1968. A comprehensive approach to human factors in developing countries. *Human Factors* 10, 6 (1968), 649–662.
- [29] Margaret S Livingstone, Warren W Pettine, Krishna Srihasam, Brandon Moore, Istvan A Morocz, and Daeyeol Lee. 2014. Symbol addition by monkeys provides evidence for normalized quantity coding. *Proceedings of the National Academy of Sciences* 111, 18 (2014), 6822–6827.
- [30] Edwin A Locke. 1968. Toward a theory of task motivation and incentives. *Organizational behavior and human performance* 3, 2 (1968), 157–189.
- [31] Microsoft. 2018. Azure Functions–Serverless Architecture. <https://azure.microsoft.com/en-us/services/functions/>. (2018). (Accessed on 02/28/2018).
- [32] Microsoft. 2018. Azure Machine Learning Anomaly Detection API | Microsoft Docs. <https://docs.microsoft.com/en-us/azure/machine-learning/team-data-science-process/apps-anomaly-detection-api>. (2018). (Accessed on 09/30/2018).
- [33] Microsoft. 2018. Working with the App Service Mobile Apps managed client library. <https://docs.microsoft.com/en-us/azure/app-service-mobile/app-service-mobile-dotnet-how-to-use-client-library>. (2018). (Accessed on 03/02/2018).
- [34] Phil Nash. 2018. A set of useful Twilio Functions. <https://github.com/phlnash/useful-twilio-functions>. (2018). (Accessed on 03/01/2018).
- [35] Nodemailer. 2018. Nodemailer. <https://nodemailer.com/about/>. (2018). (Accessed on 03/01/2018).
- [36] Peoples Bank of China. 2018. Peoples Bank of China Issues Barcode and QR Code Payment Rules. <http://www.pbc.gov.cn/english/130721/3456052/index.html>. (2018). (Accessed on 03/02/2018).
- [37] Department of Health and US Government Human Services, Office of Inspector General. 2000. Recruiting Human Subjects: Sample Guidelines for Practice (OEI-01-97-00196; 6/00). <https://oig.hhs.gov/oei/reports/oei-01-97-00196.pdf>. (2000). (Accessed on 03/02/2018).
- [38] Central Bank of Kenya. 2005. Kenya Electronic Payments and Settlement System (KEPSS) Rules and Procedures. <https://www.centralbank.go.ke/wp-content/uploads/2016/08/KepssRules1.pdf>. (2005). (Accessed on 03/02/2018).
- [39] University of Massachusetts Research Administration and Compliance. 2018. IRB Guidelines | Research and Engagement | UMass Amherst. <https://www.umass.edu/research/compliance/human-subjects-irb/guidance>. (2018). (Accessed on 03/02/2018).
- [40] OpenDataKit. 2018. Open Data Kit - Deployments. <https://opendatakit.org/about/deployments/>. (2018). (Accessed on 03/01/2018).
- [41] OpenMapKit. 2018. OpenMapKit Website. <http://openmapkit.org/>. (2018). (Accessed on 03/01/2018).
- [42] Steven Ovadia. 2014. Automate the internet with “if this then that” (IFTTT). *Behavioral & social sciences librarian* 33, 4 (2014), 208–211.
- [43] Particle. 2018. Particle Guides | Webhooks. <https://docs.particle.io/guide/tools-and-features/webhooks/>. (2018). (Accessed on 03/02/2018).
- [44] PayPal. 2018. Creating and managing NVP/SOAP API credentials. <https://developer.paypal.com/docs/classic/api/apiCredentials/?mark=API%20security>. (2018). (Accessed on 03/02/2018).
- [45] Irene M Pepperberg and Ken Nakayama. 2016. Robust representation of shape in a Grey parrot (*Psittacus erithacus*). *Cognition* 153 (2016), 146–160.
- [46] petnet. 2018. Petnet.Io Smartfeeder - The SmartFeeder feeds your pet the right amount at the right time, automatically! <https://www.welcome.ai/products/hardware-iot/petnet-io-smartfeeder>. (2018). (Accessed on 09/30/2018).
- [47] Maria Elena Miletto Petrazzini. 2014. Trained quantity abilities in horses (*Equus caballus*): A preliminary investigation. *Behavioral Sciences* 4, 3 (2014), 213–225.
- [48] Arjuna Sathiaselalan, Richard Mortier, Murray Goulden, Christian Greiffenhagen, Milena Radenkovic, Jon Crowcroft, and Derek McAuley. 2014. A feasibility study of an in-the-wild experimental public access wifi network. In *Proceedings of the Fifth ACM Symposium on Computing for Development*. ACM, 33–42.
- [49] Peter D Schellie. 1979. Electronic Fund Transfer Act. *The Business Lawyer* (1979), 1441–1452.
- [50] Eleanor Singer and Cong Ye. 2013. The use and effects of incentives in surveys. *The ANNALS of the American Academy of Political and Social Science* 645, 1 (2013), 112–141.
- [51] Square. 2018. Square Connect API Documentation. <https://docs.connect.squareup.com/api/oauth#credentials>. (2018). (Accessed on 03/02/2018).
- [52] RG Stovel, S Ginsburg, L Stroud, RB Cavalcanti, and LA Devine. 2018. Incentives for recruiting trainee participants in medical education research. *Medical teacher* 40, 2 (2018), 181–187.
- [53] Stripe. 2018. Stripe API Reference. <https://stripe.com/docs/api#authentication>. (2018). (Accessed on 03/02/2018).
- [54] SurveyCTO. 2018. Homepage. <https://www.surveyccto.com/>. (2018). (Accessed on 07/22/2018).
- [55] R Core Team et al. 2013. R: A language and environment for statistical computing. (2013).
- [56] Twilio. 2018. Twilio - Communication APIs for SMS, Voice, Video and Authentication. <https://www.twilio.com/>. (2018). (Accessed on 03/02/2018).
- [57] Twilio. 2018. Two Factor Authentication for Identity Management via SMS or Voice. <https://www.twilio.com/use-cases/two-factor-authentication>. (2018). (Accessed on 09/30/2018).
- [58] Dale Whittington, John Briscoe, Xinming Mu, and William Barron. 1990. Estimating the willingness to pay for water services in developing countries: A case study of the use of contingent valuation surveys in southern Haiti. *Economic development and cultural change* 38, 2 (1990), 293–311.
- [59] Xinglin Zhang, Zheng Yang, Wei Sun, Yunhao Liu, Shaohua Tang, Kai Xing, and Xufei Mao. 2016. Incentives for mobile crowd sensing: A survey. *IEEE Communications Surveys & Tutorials* (2016).
- [60] Xinglin Zhang, Zheng Yang, Zimu Zhou, Haibin Cai, Lei Chen, and Xiangyang Li. 2014. Free market of crowdsourcing: Incentive mechanism design for mobile sensing. *IEEE transactions on parallel and distributed systems* (2014).