The Practice of Informatics JAMIA

Viewpoint Paper ■

e-Consent: The Design and Implementation of Consumer Consent Mechanisms in an Electronic Environment

ENRICO COIERA, MBBS, PHD, ROGER CLARKE, MCOMM, PHD

Abstract The effective coordination of health care relies on communication of confidential information about consumers between different health and community care services. However, consumers must be able to give or withhold "e-Consent" to those who wish to access their electronic health information. There are several possible forms for e-Consent. In the general consent model, a patient provides blanket consent for access to his or her information by an organization for all future information requests. Conversely, general denial explicitly denies consent for information to be used in future circumstances, and in each new episode of care, a new consent would be needed to obtain information. In the general consent with specific denial model, a patient attaches specific exclusion conditions to his or her general approval to future accesses. In contrast, in the general denial with explicit consent model, a patient issues a blanket block on all future accesses but allows the inclusion of future use under specified conditions. There also are several alternative functions for an e-Consent system. Consent could be captured as a matter of legal record. E-Consent systems could be more active by prompting clinicians to indicate that they have noted consent conditions before they access a record. Finally, the record of patient consent could be fully active and used as a gatekeeper in a distributed information environment. There probably will need to be some form of data object that is associated with patient information. This e-Consent object (or e-Co) will contain the specific conditions under which the data to which it is attached can be retrieved. Given the complexity of clinical work and the substantial variation we can expect in an individual's desire to make his or her personal medical details available, it is unlikely a "one size fits all" approach to e-Consent will work. Consequently, with a well-chosen consent design, it should be possible to balance the specific need for privacy of some of the population against the desire by others to err on the side of clinical safety, and clinicians desire to minimize the burden that an electronic consent mechanism would impose.

■ **J Am Med Inform Assoc.** 2004;11:129–140. DOI 10.1197/jamia.M1480.

The effective coordination of health care relies on the communication of confidential information about consumers between different health and community care services. Electronic data exchange and Internet technologies increasingly play important roles in such communications. Consumers must, however, be able to give or withhold consent to those who wish to access their electronic health information.

For example, electronic patient records are seen by many as an essential prerequisite for health care,[1] opening up patient data to the whole clinical team involved in patient care. So, by definition, the presence of an electronic environment means that more clinical workers will be able to access patient information more often and in a greater diversity of locations. With the broadening of access to patient information comes the risks that such information is used for purposes not originally consented to by the patient.

While much is known about the ways in which security technology can protect information transactions from unwanted interception,[2] very little work exists to determine how a consumer's consent to view their private information is safeguarded in a networked and online environment. This

Affiliations of the authors: Centre for Health Informatics (EC); Baker Cyberspace Law & Policy Centre, University of New South Wales (RC), Sydney, Australia.

This work was performed under government contract, and the copyright belongs to the Australian Department of Health and Ageing. The authors published this work previously in part at the government Web site ⟨http://www.health.gov.au/hsdd/primcare/it/pdf/coiera.pdf⟩. The opinions expressed in this article are not necessarily the views of the Australian Government Department of Health and Ageing.

The authors have obtained permission to publish this revised work in *JAMIA* from the Australian Department of Health and Ageing.

Thanks are due to the other members of the e-Consent team for their intellectual contributions to this report: Peter Broadhead, John Payne, Jane Spittle, and Christine O'Keefe. Full details of the e-Consent project are available at http://www.health.gov.au/hsdd/primcare/it/econsent.htm.

Correspondence and reprints: Enrico Coiera, MBBS, PhD, The Centre for Health Informatics, The University of New South Wales, UNSW, NSW, 2055, Australia; e-mail: <ewc@pobox.com>.

Received for publication: 10/20/03; accepted for publication: 11/06/03.

report will outline a framework for obtaining and determining electronic consent (e-Consent) within health care. It will examine a range of models for e-Consent and examine some of the technical issues associated with transforming those models into working systems. It is not the intention of this report to make a specific judgment about which consent models are more acceptable or to make specific recommendations about the detailed implementation of an e-Consent system. Such decisions would need to reflect the legal framework within which any e-Consent system operates, and the expressed wishes of consumers regarding the strength of protection they desire.

Specifically, the report proposes a set of basic design principles that any consent framework might need to adhere to and focuses on some of the trade-offs in system performance that these principles imply. It then examines various possible forms of consent, explores the ways that these can be implemented in an online environment, and examines how well these models reflect the design principles. Next, the report explores the nature of information exchanges in health care and uses this to reflect on the acceptability of the different consent mechanisms in the clinical workplace, as well as to consumers. Finally, the report develops a health transaction model and uses this to sketch the set of behaviors or services an e-Consent system will need to perform its key functions. Appendices contain detailed examples of information transactions in health and an example set of computational rules for determining consent.

Translating Existing Consent Processes into the Online Environment

One of the challenges to the design of an electronic mechanism is to make sure that the translation of legal rules designed to regulate human activity does not have unexpected consequences when implemented in an electronic environment.

Currently, individuals working in the health system usually are responsible for obtaining consent from a patient or determining whether consent exists, prior to accessing, using, or passing on a patient's information. In an electronic environment, the existence of patient consent may be determined by automatic processes without the explicit involvement of the parties normally associated with that decision. For example, clinical staff working in a hospital might have their right to access electronic patient records determined by a set of computer rules that attempt to assess whether they have reasonable reason to access any particular data.

The effects of implementing a legal framework for consent in a human decision environment may be far different from that of its implementation in an information system. For example, the rules governing consent to view health information are captured in legislative frameworks, such as the Health Insurance Portability and Accountability Act (HIPAA) framework in the United States.[3] However, such rules are usually designed to govern the activities of people, not an information system. In an electronic environment, legal rules that are meant to be used by humans become rigidly codified in computer behavior. This may have the unintended consequence of increasing the number of consent actions that clinical staff undertake, increasing their workload.

There probably is a "consent gap" between the strict letter of legal requirement and current practice. An e-Consent system may enforce a closing of the gap so that information usage more closely tracks legislative intent. Closing the gap will introduce new work for clinicians, and independent of the appropriateness of that intervention, will affect clinical work. Historically, when information systems introduce unwelcome burdens on users, they are poorly accepted. Consequently, along with the introduction of an e-Consent system, there will need to be a process of user education, culture change, and, in overworked sectors, effort needs to be made to minimize the impact of imposing "new" work by delivering other economies.

Another unintended consequence of the translation of consent rules to the electronic environment is that access to patient data may become harder. If, for example, a patient is in a life-threatening state, in the paper world, as long as the paper patient record is physically present, then it can potentially be accessed by anyone present. However, in an online environment, information access may be prevented, because clinicians do not have explicit access rights. Thus, when patient consent is strongly enforced, patient safety may be compromised. Conversely, without the existence of some e-Consent mechanism, the widespread distribution of patient information across distributed networks would be open to widespread access by individuals who were not given consent, substantially breaching personal privacy.

e-Consent Design Principles

To ensure that an e-Consent system both exhibits the behaviors we expect of it and minimizes the likelihood of unwanted or unexpected behaviors, we need to develop a set of design principles. Any proposal for a system design or specific technologic solutions should then be judged against these design principles. This permits us to discuss the behavior we expect of computer-mediated patient consent, independent of complex arguments about underlying technology infrastructure.

The following set of principles is proposed as a basis for discussion. An electronic consent system:

1. Permits access to confidential patient information by checking that patient consent exists for the information request by invoking methods that check for explicit, inferred, or implied consent.
2. Should allow access to patient information to those individuals who have been explicitly permitted by a patient to view their information.
3. Should never allow access to patient information to individuals who have been explicitly denied access by a patient.
4. Should allow access to patient information to individuals who can be determined to have inferred or implied consent on the basis of their clinical role or responsibility or the clinical circumstance.
5. Does not endanger patient safety by denying access to information by clinically approved individuals where consent is either indeterminate or in defined circumstances denied.
6. Does not impede clinical work by denying access to information by clinically approved individuals, where consent is indeterminate.

7. Has security safeguards that prevent provably unauthorized individuals from accessing patient information by circumventing the consent checking mechanism.
8. Minimizes the number of requests it makes to clinicians and patients to avoid unnecessary impediment or disruption of the clinical process or the private lives of individuals.
9. Does not require an expensive or burdensome administrative infrastructure to support the obtaining and determining of consent and performance monitoring of the system.

Several conflicts exist in these principles and, as a result, a set of trade-offs results, which need to be balanced by gauging public desire, legislative frameworks, and system design. For example, there is an explicit trade-off between privacy and clinical safety that humans make routinely that will potentially be substantially altered in a computer environment. Principles 1 through 4 are in conflict with Principles 5 and 6. The latter are designed to err on the side of clinical judgment to maximize patient safety. The former are designed to rigidly enact the legal framework. Similarly, Principles 8 and 9 are designed to minimize the impact of any consent checking system, which could intrude into almost every request for patient information. Again, this design requirement is in conflict with Principles 1 through 4, which ask for rigid checking of consent before any information is passed across.

In the next two sections, we explore several different models for patient consent and their electronic implementation and see how well these models perform in relation to the consent principles above. As will become apparent, different models perform well with some principles and poorly against others, and they strike different balances in the trade-off between privacy, safety, and effect on clinical work practices.

Forms of Patient Consent

In an effort to avoid the potential complexity of consent systems, some people advocate the use of a blanket "opt-in" or "opt-out" model. In the "opt-in" model, a patient essentially provides blanket consent for access to his or her information to an organization for all future information requests. For example, a patient may choose to join a health organization that has an electronic patient record as a part of its service. At the time of joining the service, the patient provides consent for all information requests about his or her personal medical information by identifiable employees of the organization. Conversely, the "opt-out" model explicitly denies consent for information to be used in future circumstances, and in each new episode of care, a new consent would be needed to obtain past information. In between these two extreme models lies a range of possibilities in which consumers specify under which circumstances they will permit health workers to access their information, and for which they want to deny access.

Four different forms of consent can be identified:

1. *General consent*. This is a "blanket consent" given by a patient for any health care professional working within a specified health context to access any and all of their health information for any purpose relating to the consumer's care. It corresponds to the "opt-in" model and persists into the future, unless specifically revoked by the patient. Thus, in a future episode of care, beyond the one in which the initial consent is given, a health worker does not need to ask for consent when looking at details of past episodes of care.
2. *General consent with specific denial(s)*. In this case, a patient provides a general consent, but denies consent as follows:
   - to the disclosure of particular information, and/or
   - to the disclosure to a particular party or category of parties, and/or
   - to the disclosure for a particular purpose.

   Thus, the consent to blanket future access of information is modified by specific identified conditions in which consent is to be withheld. For example, a patient provides a general consent to a health care professional but expressly precludes the disclosure of information about a sexually transmitted disease (STD) condition or gynecologic procedure, disclosure to his or her immediate family or family general practitioner (GP), or disclosure for purposes other than the treatment of a heart condition.
3. *General denial with specific consent(s)*. In this case, a patient denies all access to his or her health data, except for circumstances that are the subject of specific consent:
   - to the disclosure of particular information, and/or
   - to a particular party or category of parties, and/or
   - to the disclosure for a particular purpose.

   For example, a patient authorizes his or her primary treating professional to provide a sample of body fluids, accompanied only by relevant data, to a diagnostic service, for the purpose of conducting a battery of tests and reporting the results back to the GP. The use of patient data for population health research is another example of a specified use. In many countries, population health uses are specifically covered by legislation and are therefore beyond the legal scope of patient consent, but could easily be incorporated into the consent structure if there were scope for personal choice.
4. *General denial*. In this case, a patient provides a "blanket denial" and expressly denies consent for information to be used in future circumstances. Here, in each new episode of care, a new consent would need to be obtained. This is equivalent to the opt-out model. Contexts in which consumers would be likely to use this are treatments for STDs, drug rehabilitation, and psychiatric treatment. General denial is essentially a specified-purpose model, in which a patient is requested for consent to access information each time there is a new request by a clinician.

With general consent with specific denial, and general denial with specific consent, every request for information in the future is tested against some set of rules. We would expect with a general consent with specific denial that information access requests generally would be approved, but in the general denial with specific consent, most requests would be denied. If a patient's wishes are fully effected, each form of consent trades off different degrees of information protection against ease of access, and there is a continuum of protection offered by the different forms of consent (Fig. 1).

While some forms of consent offer patients precise control over access to their personal information, not all patients may use it effectively. More generally, all consent decisions are complicated by the fact that many patients will lack an

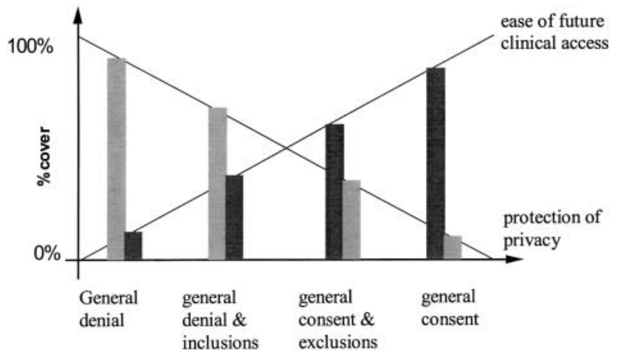

Figure 1. Different forms of consent balance clinical access and patient privacy in different proportions.

understanding of the true effects of their consent decisions. For example, patients often will be unaware of the complex relationships between health care entities, and the secondary uses to which their records might be put, or fail to recognize the implications of their consent decisions in unforeseen scenarios of their own care.

Functions for e-Consent

When modeling consent, we need to separate out two distinct concepts:

- The form of the consent—what does the record of consent look like?
- The function played by the consent—what is the operational impact of the consent?

Having explored various forms of consent, the next question to address is the effect that we may wish the form of consent to have in an electronic environment. How fully do we wish to implement a patient's wishes into the behavior of the e-Consent system? There are essentially three broad functions to consider:

1. *Patient consent can be captured as a matter of legal record only.* The burden of adhering to the patient's wishes rests with clinicians and workers in the health system, as is the case in the nonelectronic environment. Thus, the electronic record of a patient's consent is passive to the information transactions of patient data once consent has been recorded but can be called upon if there is a dispute, to bear witness to the patient's signified intentions.
   An e-Consent system of this form may still be quite active in capturing the consent record. Using structured data entry, clinicians could be guided to capture the specific form of consent granted, including inclusion and exclusion criteria from predefined lists where possible. In its most prescriptive incarnation, clinicians using an e-Consent record system cannot proceed with actions on the computer system until the consent process has been documented.
2. *An e-Consent system can actively require clinicians to signify they understand specified consent prior to accessing information.* Here, the e-Consent system takes a more active role, accessing the record of patient consent, providing it to those about to access information, and checking that they understand the terms of the consent. Specifically, it could prompt individuals wishing to access a record that the patient has expressed particular wishes about the use of their information and that a record of those wishes is available. The individual wishing to access information may then be asked to acknowledge his or her understanding of that fact, or that he or she has accessed the consent record and understood it, before accessing the patient record. However, the decision to access information remains with the clinician, and there is no bar to access patient information. This type of system can generate an audit trail of accesses to information and can be used to both retrospectively check that patient consent is being observed or, in cases of dispute, can act as an authoritative record.
3. *The e-Consent system could act as a gatekeeper and permit only consented individuals to access information*. This is the most active class of behavior possible. In a distributed information environment, many individuals may choose to access patient information. If we wish to ensure only those with proper authorization view patient data, then the e-Consent system checks to see that the individual who wishes to access the information is able to satisfy the conditions of consent before access is granted. There would thus need to be a "consent machine" implemented that can read the consent record and match the conditions in the record with the individual seeking to access information. To conform with the design principles outlined earlier, a gatekeeper system would need to have a method to allow it to be overridden, for example, when a clinician is denied access but believes the situation is a medical emergency.

As with the four models of consent identified in the previous section, each of these three functions has a distinct potential for impact on clinical work and patient privacy (Fig. 2). Clinical work and patient safety are least affected by the use of the consent as a legal record only. Clinical work is slightly more disrupted when the requirement to signify consent is understood before accessing patient records and most affected by the e-Consent system acting as an active gatekeeper. Patient privacy, on the other hand, is least protected by treating consent to be a matter of record only, is more protected by the generation of an audit trail of acknowledgment of consent, and is most protected by the gatekeeper function.

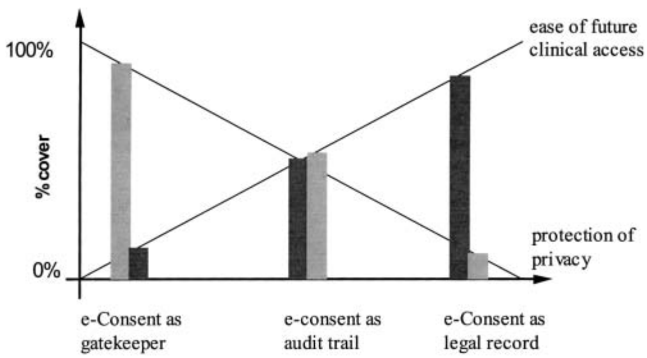

Figure 2. The different possible functions of consent balance clinical access or patient privacy in different proportions. The diagram is illustrative of the balances only; thus, there is no intention to portray the balance between access and privacy as equal in the middle model of e-Consent as an audit trail.

Table 1 ▪ Relative Ranked Effect of Different Combinations of Consent Form and Function on Protection of Patient Privacy (P1 = minimum, P12 = maximum) and Ease of Clinical Information Access (C1 = minimum, C12 = maximum)

Form Effect	General Denial	General Denial, Explicit Consent	General Consent, Explicit Denial	General Consent
Record	P4 C9	P3 C10	P2 C11	P1 C12
Audit	P8 C5	P7 C6	P6 C7	P5 C8
Gatekeeper	P12 C1	P11 C2	P10 C3	P9 C4

Thus, there are two design dimensions available to us to first capture and then enact consent—the form of consent axis and the effect of consent axis. The decision to implement an e-Consent system as a passive record, active prompt and audit trail, or gatekeeper will probably have the largest impact on the privacy and clinical impact trade-off. The four forms of consent offer a degree of fine-tuning by permitting patients to choose which one of the forms of consent they would like. If patients are given appropriate information, they should be able to make informed choices about the extent of "consent cover" they wish to select, recognizing the degree of clinical risk that comes with that choice.

Table 1 combines the orderings in Figures 1 and 2 to give a combined relative ordering of the impact on privacy and clinical work of different combinations of forms of consent and consent function. This creates a relatively rich space of designs for an e-Consent system. Although there are clear conflicts within the original design principles, there thus is a broad range of options available to select a solution that closely satisfies the privacy versus clinical work trade-off considered most suitable for a given setting.

In a technical design sense, the main decision will thus be to choose one of the three functions that defines the effect of consent on clinical work. The forms of consent only constrain the set of specific circumstances that can be expressed within a given function. For example, independently of the form consent chosen, a gatekeeper system will still need to check whether a request to access information satisfies a specific set of conditions. That will require the construction of a "consent engine" capable of performing such condition checks, but its core design will not be altered if the conditions are the inclusion of specific denial or consent conditions (see section on A Computational Model of e-Consent). There thus appears to be no technical need to build to a specific function and be limited to only one of the four forms of consent. Rather, we should be able to accommodate all of the four forms of consent in any of the functions that are chosen.

The Impact of e-Consent on Clinical Work

While much public debate centers on patient privacy and the need to protect patient information at all costs, little is said about the nature of clinical work and the likely impact that introducing a more restrictive access process to information would have on patient care. Currently, our detailed understanding of information transactions in the health care system is sadly limited. Consequently, it will be difficult to make precise judgements about the effect of introducing any of the previous consent models into the clinical workplace. However, we are in a position to make some general observations about the likely impact of e-Consent systems on clinical work.

Box 1—Information Transactions in Clinical Work

There are only sparse data available to describe the volume of information transactions in health care, especially comparing formal information system transactions with informal interpersonal ones. Covell and colleagues reported that about 50% of information requests by clinicians in a clinic were met by a colleague rather than documented sources.[4] Tang et al. found that about 60% of a clinician's time in a clinic is devoted to talk.[5] Even in a hospital with a mature computerized information system, 50% of the information transactions occurred face-to-face between colleagues, with e-mail and voice mail accounting for about another quarter of the total (Safran et al.).[6]

Studies conducted in emergency departments provide a rich picture of the communication transactions that need to occur to provide this health service.[7] In those data, about 82% of all information exchanges observed in accident and emergency (A&E) departments occurred as face-to-face conversations. Information exchanges involving patient notes, computer access to patient results, or paper-based forms accounted for only about 10% of information exchanges. Further, while it would be expected that in an information-rich environment, most information requests would be to obtain specific patient data such as test results, most information exchanges were associated with asking for or receiving information between colleagues.

This picture of clinical work will clearly vary from setting to setting, but as the Australian and U.S. results cited above suggest, the picture of heavy reliance on interpersonal exchanges for information transmission is repeated across most settings from about 50% in clinics to about 85% in A&E.

Clinical work is complex, and often is rapidly moving and consensual, especially in hospital environments. One of the striking features of clinical work is that clinicians rely far more on informal interpersonal information exchange than formal paper or computer records to satisfy their information requests (Box 1). This is in part because information is not always easily available from formal sources, but more often is simply because conversation is a more appropriate mechanism for information exchange. Thus, formal information systems such as the electronic patient record are a necessary, but not sufficient, source of information to support clinical work.

With this in mind, there are a number of consequences for any e-Consent system.

- First, e-Consent systems that regulate access to electronic patient information will protect only a small proportion of all information transactions in the health system. Most patient information will be exchanged through conversation, whether face-to-face, by e-mail, or on the telephone. The current clinical processes and legislative framework regulating patient consent will need to continue to cover these interactions.
- Second, with time, conversations will start to come under the electronic umbrella. Most voice telephony services already transmit data in digital form, and, in the future, conversations are as likely to be archived on computer as e-mail is today. Consequently, it should be feasible at some point to automate searches through conversational databases to retrieve confidential information. E-Consent will need to evolve to mange this form of record since it contains patient data within it.

Box 2—Summary of Information Transaction Attributes

The basic model for a transaction can be summarized as follows:

A transaction is an event involving:
1. The transmission of data
2. About a patient agent
3. For a defined purpose
4. By an initiating agent
5. To one or more receiving agents (which may be people, storage, or programs)
6. Across a communication channel
7. Using a communication service
8. Accessed via a communication device

An agent has:
1. An identity
2. A set of roles (which may include patient and clinical roles)
3. Access to one or more communication channels
   a. Via one or more communication services
   b. Via one or more communication devices

The transaction is consented:
1. Explicitly by the patient or
2. Implicitly by the application of rules to the transaction description

Transaction rules determine whether consent can be logically inferred for a proposed transaction from the existing consent explicitly given by a patient or from the status of the requesting and receiving agents as they relate to the proposed purpose.

- Embedding routine consent checks on every request for information through an electronic patient system (whether as a formal gatekeeper function or to record clinician understanding) in an environment such as accident and emergency (A&E) may substantially impede clinical work. Clinicians often are working to capacity and juggling multiple tasks. Imposing extra duties because of the introduction of an information system, without removing existing tasks, will limit the capacity of clinicians to perform their jobs in a time-critical way.

Consequently, the emphasis in behavior we would expect of an e-Consent system needs to vary with the clinical setting in which it is used. Depending on the clinical transaction load and the need for patient confidentiality, there will be variations in the balance between permissive and prescriptive access to patient information for clinicians. It is likely that we cannot second-guess which combination of form of consent and function ideally matches the needs of a given clinical environment and that some experimentation is needed to gather data about the acceptability of any given model.

The Basic Transaction Model

So far we have described a set of models for the form and function of e-Consent. Before a detailed description of the design of an e-Consent system can be developed, we will need to develop a general model of information transactions that describes the communication events that could occur in a given clinical setting and that identify which consent mechanisms are needed (Box 2).

In the basic model, we define a communication event as a transaction involving the exchange of information between two agents to satisfy a particular health goal. For the transaction to occur, a communication channel will be selected and a message sent across that channel. Typically, the agent will communicate via a device, and one of a number of communication services will be available on that device.

For example, a GP may call a specialist to discuss a patient. In this case:

- the agents are the GP and specialist
- the device is the telephone
- the channel is the public telephone network
- the service is voice telephony

If the specialist were unavailable, the GP might use an alternative service such as voice mail, or leave a message via an alternate agent like the receptionist at the specialist's rooms. Other channels that might be available include the Internet with services such as e-mail. Thus, the model also needs to include channels such as the Internet and infrared transmission between information devices, communication services such as e-mail and voice-mail, and devices such as mobile phones and palm-top computers. It also needs to recognize that sometimes one of the agents may be a computer. For example, the GP in our example may request laboratory test results about a patient by logging on to a database.

We next note that the messages between agents contain information that may relate to an individual patient. The patient will have consented to give the information for a particular purpose, with consequent consent for that information to be shared, transmitted, or stored in ways essential to satisfying that purpose. Sometimes that consent is explicitly asked for and sometimes it may be inferred as existing based on a set of consent assignment rules (e.g., in the gatekeeper model).

As a consequence, we need to ensure that any data required to determine whether an assignment rule is satisfiable will also be captured in the model. For example, we may need to model the identity and roles that the agents have in satisfying a given purpose, since these have a direct bearing on consent. In the GP example above, the purpose may be to get advice on the management of the patient's case or discuss whether referral is appropriate. To satisfy the task, confidential patient information is exchanged. The role of the GP is the designated primary caregiver of the patient, having explicit consent from the patient to manage his or her care and to access confidential information required to carry out that task. The specialist's role is to advise the GP on the best treatment for the patient. Consent for the specialist to be exposed to the confidential patient information is implied only, since the patient was not consulted prior to making the call.

For example, rules to infer that consent exists might be of the following form:

Infer consent for proposed transaction X if

> Patient has consented explicitly to purpose A and
> Transaction X is an essential component of the task set needed to satisfy purpose A.

Infer consent for proposed transaction X if

> Patient has consented to agent A to act on his or her behalf in relation to purpose B and
> Agent A consents to transaction X.

Do not infer consent for proposed transaction X if

Purpose of X requires security level A and
Channel is security level B and
B is less secure than A.

The actual rules that need to be created will emerge from a technical reading of appropriate legislation, consultation with clinicians and consumers, and from technical requirements associated with engineering the transaction model and the inference mechanisms. The effect of implementing such rules is to decide who has a "need to know" and who has "consent to know" specific information about a patient's health. The degree to which one develops such rules depends strongly upon the function chosen. The complexity of rules needed to perform a gatekeeper function would, for example, be much greater than in the legal record model of e-Consent. Appendix 2 develops an example of a basic set of consent rules, and Appendix 3 applies them to some typical clinical transactions. Appendix 4, Horn clause logic, is discussed in Appendix 2. (Appendices 2–4 are available as online data supplements at www.jamia.org.)

If such rules are to be developed, then effort will need to be made to fully understand the dimensions of consent in the basic transaction model, with specific attention to defining:

- categories of information
- categories of entity
- categories of purpose

Even apparently simple transactions of patient information, when analyzed according to this model, reveal substantial complexity. Appendix 1 takes several clinical examples of information transfer, and decomposes them according to the transaction model developed here. It quickly becomes apparent that trying to develop a detailed model of any information transaction involving some form of coordinated care is a large task. One would also need to update such detailed models routinely as health services were restructured or as new communication technologies were adopted.

Basic Consent Services

Using the transaction model, it is now possible to describe a set of basic services or discrete functions that should be provided by an e-Consent system. Not all of these services would be used in all clinical contexts. For example, in some settings, it may be too restrictive to determine consent on a case-by-case basis during a hospital admission, and membership of an organization is all that is checked by the e-Consent system before access is allowed. The specifics of how these services would be implemented in different environments still need to be explored.

As an initial proposal, an e-Consent system may need to provide the following services:

1. Check an individual's identity, whether patient or health worker—we need to be able to verify that the clinicians requesting information are who they say they are or verify the identity of a patient giving a consent instruction.*
2. Check that an individual is employed by a health organization—when consent is assigned on the basis of membership of an organization.
3. Check an individual's clinical roles within an organization—when consent is assigned on the basis of role within an organization.
4. Check that a health organization exists.
5. Recognize a set of defined purposes for which consent may be requested—when consent is assigned on the basis of specific use of information.
6. Record the fact that consent has been given or denied for accessing clinical data by a patient or authorized agent.
7. Retrieve a consent instruction associated with a particular clinical datum.
8. Associate consent at the level of specific data, data bound by an episode of care, or data bound to a whole patient record—allowing consent to be determined at any one of these levels.
9. Be capable of recording a complex set of consent instructions, with inclusions and exclusions to access information.
10. Infer the existence of consent on the basis of a requesting individual's identity, role in an organization, or association with an organization.
11. Record the delegation of consent by a clinician to another clinician, role, or institution.

Security Services

An e-Consent system needs to be supported by a set of security functions that minimizes the likelihood of unauthorized access to patient data and that decisions are made about consent based on the best quality data.

The security layer will provide a set of services that would be made available to the consent system, including the authentication of data, integrity of data, nonrepudiation, and, where necessary, secrecy. For example, the fact that a patient has given or withheld consent should be nonrepudiable. Equally, the consent system would have the benefit of knowing that data to which it has access are authentic.

Thus, every time a fact is entered into the database (using the "assert" action in the examples in the appendices), it should probably have a degree of authentication, the level of which will vary with the data being entered. Thus, asserting a fact that a patient has given consent to release data for view would invoke an authentication service that ensured that the system is confident the patient and the clinician making the statement were indeed who they said they were and perhaps that consent was informed.

These security services are assumed to be available to an e-Consent system for the purposes of this report. The specific way in which, for example, identity is verified is an issue beyond the scope of consent. However, we need to recognize that the consent and security services will interact with each other and call on each other to perform tasks the other depends upon.

Designing and Implementing an e-Consent System to Fit within the Clinical Environment

As we have seen, the form of consent and its function both shape the impact of an e-Consent system on clinical work. A further dimension that will shape impact on clinical work is

*More precisely, we need only to check that the person giving consent is the same person to whom the data relate. A true identity is not needed, only an identity label that persists in the system. A consumer may elect to have multiple identities to prevent data cross-checking.

Box 3—Example Use of Scenario One

For a given clinical event, a clinician needs to enter patient's data into a software system such as a prescribing package or an electronic record.

- At the beginning or end of that event, the clinician would be prompted to obtain the patient's formal consent for the intended uses of that information.
- As default, the e-Consent system could operate on a general consent basis, unless told otherwise by the clinician.
- A simple menu, presenting the different forms of consent, would be offered for selection if that default behavior were not acceptable to the patient. At this time, the clinician could make a selection of an alternative form of consent if requested.
- Selection of a general denial, for example, would signify that the patient data should not be accessible beyond the clinician who is recording it and for no other purpose than the immediate episode of care.
- If the patient has specific conditions he or she wishes to stipulate, then these are captured at this moment. These conditions would include the denial or enabling of access to information to named individuals, institutions, roles, or specified uses. As far as possible, these conditions would be drawn from a predetermined list and would be offered as part of a precompiled menu of options.
- Once discussed, the clinician would signify the occurrence of the conversation, which could be automatically logged and verified, perhaps with a digital signature, insertion of an identifying card, or some password mechanism.

the specific design and implementation of the software system. It is conceivable, for example, that a poorly designed e-Consent system that only captures consent as a matter of record is more intrusive in a clinical environment than a well-designed gatekeeper system. Thus, independent of any given e-Consent model we chose, we would want to minimize the number of times a clinician had to enter data and minimize the cognitive load imposed on clinicians.

Thus, the way a consent form or function is actually implemented can have a substantial impact on clinical work. The appropriate choice of default forms of consent should minimize the impact of a system in specific environments. For example, a family planning clinic may recognize the routine need for extra protection for their patients and choose general denial as their default form of consent. Where wider clinical accessibility is desired, the general consent option is likely to be the mainstay. For example, general consent would probably be the default for a patient being admitted to a hospital, making all the patient's information available to clinical staff without further interaction between the patient and the consent system. Access in this case would be controlled by the security component of the system, which ensured individuals requesting information access were authorized members of staff.

Compare, for example, the following two use scenarios for obtaining consent. In the first example (Box 3), a system captures patient consent as a matter of record by requiring a clinician to enter the form of consent and attaching it to the patient record. In the second example (Box 4), the e-Consent system acts as a gatekeeper and is driven by the use of smart cards. In terms of impact on clinical work at the point of data entry, the smart card model is much less intrusive. On the other hand, the gatekeeper may have downstream consequences on clinical work if others attempt to access data and are denied it.

Box 4—Example Use of Scenario Two

For a given clinical event, a clinician needs to enter patient's data into a software system such as a prescribing package or an electronic record.

- At the beginning of that event, the clinician and patient are prompted to insert their smart cards into the system.
- The system automatically notes the identity of patient and clinician.
- The system automatically notes the patient's default form of consent.
- The clinician asks the patient if he or she wishes to alter the form of consent displayed on the screen.
- If needed, the consent is altered for the current set of data to be captured. In most cases, we assume that the default form of consent on the patient card will suffice.

Another design choice to be made is in the number of consent modifiers that are made available to patients to capture specific inclusion and exclusion criteria. While enhancing consumer choice, the inclusion of a large space for choice about consent options through the modified models generates some system design and maintenance complexity and costs. However, without altering the basic rights to consent of a patient or the options they may have, we can minimize the complexity of options health workers would need to navigate in routine cases. For example, the range and type of exclusions available can be ordered to first reflect the common situations in a particular practice, and the labels used could reflect local language.

A Computational Model of e-Consent—The e-Co

We are now in a position to sketch out some of the information architecture needed to support an e-Consent system. Specifically, there will probably need to be some form of data object that records a specific consent and is associated in some way with patient information. This e-Consent object (or e-Co) will contain the specific conditions under which the data it is attached to can be retrieved.

The e-Co will contain a set of conditions identifying when the data can be viewed. These conditions can be expressed in terms of named individuals, institutions, institutional roles, and specified purposes. An e-Co thus needs to record the following form of assertion:

> Access to ⟨information⟩
> by an ⟨entity⟩
> for a ⟨purpose⟩
> in a ⟨context⟩
> is {consented to | denied}

There are many different computational representations one could choose for this information, with different capacities to represent complex conditions. At the most expressive, one could use some representation compatible with first order logic, which would allow extremely rich combinations of conditions and nesting of conditions to be captured.

A simpler list representation is possible if we elect to only have a set of basic conditions that can simply be included or excluded and not allow complex combinations. Here, an in-list contains the set of conditions that are explicitly permitted, and an out-list contains the set of conditions under which patient information specifically cannot be viewed.

Using the four forms of consent presented earlier, the list representation of an e-Co would have the following form:

General consent:

[in list—all workers covered by scope of consent]
[out list—nil]

General consent with exclusions:

[in list—all workers covered by scope of consent]
[out list—specified exclusions to override in list]

General denial:

[in list—nil]
[out list—all]

General denial with inclusion:

[in list—specified inclusions to override out list]
[out list—all]

The precise form of the e-Co and its supporting infrastructure will depend on the function for consent chosen.

- If the e-Co is simply a record of a patient's consent, then it does not need to have a complicated support structure. At its simplest, the e-Co as legal record is a free text record, and at its most complex, it is created from a predefined set of primitives that describe the common agreed elements to consent. In the latter case, there would thus need to be a database of allowed entities (name or role or institution), purposes, and contexts.
- If the e-Co is part of a more active system that prompts users to indicate they have read and understood the record, the e-Co may need to include the people who have read and asserted they understood it and when that happened. Such assertions would need to be protected by a number of security services, for example, ensuring such assertions are nonrepudiable.
- When the e-Co is part of an active gatekeeper system, not only does there have to be a database of allowed entities, purposes, and contexts from which to construct it, there needs to be a library of rules for determining whether the conditions specified in the e-Co are satisfiable by the person requesting access. This requires some form of rule knowledge base and an inference engine to interpret the rules according to the conditions specified.
- We can conceive of an even more active implementation of an e-Co, which is as a software agent. Here, patient information is "wrapped" inside the e-Co and cannot be unlocked unless the e-Co is satisfied that conditions for viewing are met. This would allow the e-Co to move widely across a distributed information system with patient data and means that individual software systems do not need to worry about implementing e-Consent mechanisms themselves.

The scope of information for which an e-Co is valid is an issue for discussion. An e-Co could be attached to (moving from most specific to most general coverage):

- an individual datum like a laboratory result.
- all the data captured in a specified episode of care.
- all the data associated with a patient identity (perhaps across several episodes of care).
- all the data associated with a patient.

Consent would be attached at the most general level appropriate. Thus, if it were burdensome to require clinicians to explicitly certify consent for every piece of clinical data, and patients had no objections, it would be ideal to obtain a certificate for a whole episode of care or for all of a patient's records.

The temporal scope of an e-Co also is an issue. In a general denial situation, we are essentially asking for and obtaining certificates for each information request. However, in situations of general or modified general consent, certificates persist over time and permit future and as yet unidentified access to information. In these cases, the e-Co would need to be modifiable by the patient after it was issued.

Conclusion

This report has outlined several possible models for determining that patient consent exists prior to allowing access to health information. Through the discussion, it has emerged that a variety of different consent behaviors are possible and that their desirability varies with health sector and individual patient preference. From a technical viewpoint, most consent models will require some mechanism for capturing specific inclusion and exclusion criteria that define a patient's explicit consent intentions. Given the complexity of clinical work, and the variation in consumer needs for consent, there is no obvious single set of design criteria that can be uniformly adopted. The specific implementation of an e-Consent system will always trade-off issues such as a patient's desire to protect confidentiality, the impact of consent systems on clinical work, and the cost of designing and maintaining a potentially highly complex system. As a result, a general approach to e-Consent is needed, which can be customized to the local needs of differing health sectors, and accommodating a variety of patient wishes.

References ■

1. Dick RS, Steen EB, Detmer DE, (eds). The Computer-Based Patient Record—An Essential Technology for Health Care. Washington DC: National Academy Press; 1997.
2. Barrows RC Jr., Clayton PD. Privacy, confidentiality: and electronic medical records. J Am Med Inform Assoc. 1996;3:139–48.
3. Tang PC. An AMIA perspective on proposed regulation of privacy of health information. J Am Med Inform Assoc. 2000;7:205–7.
4. Covell DG, Uman GC, Manning PR. Information needs in office practice: are they being met? Intern Med. 1985;103:596–9.
5. Tang P, Jaworski MA, Fellencer CA, Kreider N, LaRosa MP, Marquardt WC. Clinical information activities in diverse ambulatory care practices. Proc Am Med Inform Assoc. 1996:12–6.
6. Safran C, Sands DZ, Rind DM. Online medical records: a decade of experience. Proc EPRIMP. 1998:67–74.
7. Coiera E, Jayasuriya R, Hardy J, Bannan A, Thorpe MEC. Communication loads on clinical staff in the emergency department. Med J Aust. 2002;176:415–8.

Appendix 1 Example Transaction Models

Case 1—Doctor Requests and Receives a diagnostic Test

A general practitioner (GP) places a request for a diagnostic test in relation to a particular patient. This may involve provision of a sample of body fluid or tissue, or the presentation of the person at the diagnostic service's premises. A report is prepared. The report may be sent directly to the doctor or provided to the patient to carry to the doctor. The diagnostic service may or may not be aware of the patient's identity; may or may not retain a copy of the sample, of the report, and of materials generated during the testing process; and may or may not provide to some additional party a copy of the report or other information arising from the process.

(**Agents**:

Patient,
Referring Doctor,
[Patient nominated additional party *roles*:
patient relative/executor,
consulting clinician,
insurance company,
federal agency, e.g., immigration, etc.]
[Referring office *roles*:
practice nurse,
office clerk,
information system manager]
[Laboratory *roles*:
office staff,
laboratory staff,
clinician,
information system manager,
interpretation software]
[Courier *roles*:
office staff,
courier])

(**Data**: [A: patient ID], (*Data groupings are labeled by letter on flow diagram*)

[B: Extract of patient record:
history,
past test results,
diagnosis],
[C: Test purpose,
Test order,
Specimen],
[D: Test result,
Test interpretation and report],
[E: Referring doctor ID, address],
[F: Receiving laboratory ID, address],
[G: Courier office ID, address,
Courier record of specimen or order])

(**Data storage sites**:

[referring office storage:
paper patient record,
electronic patient record],
[laboratory office storage:
paper patient record,
electronic patient record],
[courier storage:
office records,
courier vehicle records])

Possible Transaction and Data Flows

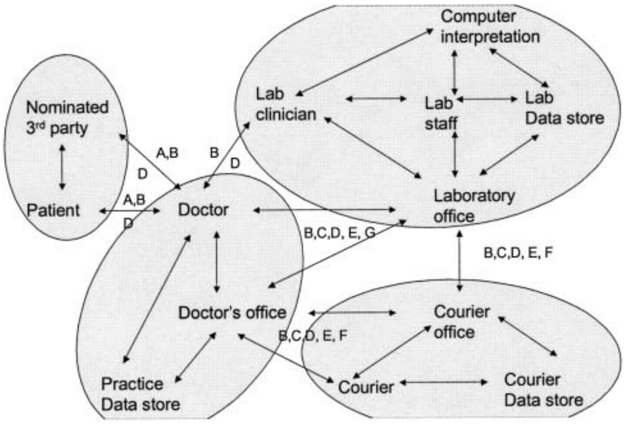

Any of the following combinations of communication channel, communication service, and communication device may be involved in transmission of some data for some of the transactions required during this scenario.

Need-to-know Matrix (Agent vs. Data to determine need to access and use specified data):

	Patient ID	Doctor ID	Lab ID	Courier ID	Purpose	Patient Record	Test Order	Specimen	Courier Record	Test result	Report
Patient	Y	Y	?		Y	Y	Y	Y		Y	Y
Referring doctor	Y	Y	Y	?	Y	Y	Y	Y	?	Y	Y
Patient-specified party	?	?	?		?	?	?	?		?	?
GP practice nurse	Y	Y	?	?	?	?	?	?	?		
GP office clerk	Y	Y	?	?	?	?	?	?	?		
GP IS manager	Y	Y	?		?	?	?	?			
Lab office staff	?	Y	Y	Y			Y		?		
Lab lab staff								Y		Y	
Lab clinician					Y	Y		Y		Y	Y
Lab IS manager	?	Y	Y								
Interpretation software					Y	Y		Y		Y	Y
Courier office		Y	Y	Y			Y	?	Y		
Courier courier		Y	Y	Y			Y	?	Y		

[Channel	**Service**	**Device]**
[Face to Face:		
	Patient (as channel)	
	Referring doctor (as channel)]	
[Physical network:		
	Post	Letter
	Courier	Packet
		Specimen]
[Wireless local network		
	IR	Palm top computer
		Printer
		Computer
		Courier bar code scanner]
[Wireless wide-area networks:		
	Paging	Pager
	Short message service	Mobile phone
	Voice	Mobile phone
	WAP	Mobile phone
		Palm-top computer]
[Telephone network		
	Voice	Telephone
	Fax	Fax machine
	Internet e-mail	Computer
		Palm-top computer
	Internet IP, e.g., Web forms	Computer
		Palm-top computer]
[Internet/intranet		
	Test order entry/	Computer
	reporting software	Printer
		Storage archive, e.g., DAT
	Internet e-mail	Computer
		Palm-top computer
	Internet IP, e.g., Web forms	Computer
		Palm-top computer]

Case 2—Discharge Summary

A Registrar prepares a discharge summary for a patient who has spent a period in the hospital. The hospital sends the document to the patient's GP.

(**Agents**:

Patient,
[Patient nominated additional party *roles:*
patient relative/executor,
consulting clinician,
insurance company,
federal agency, e.g., immigration]
[Hospital *roles*:
registrar,
specialist,
medical transcriptionist]
[GP office *roles*:
general practitioner,
practice nurse,
office clerk,
information system manager])

(**Data**: [A: patient ID], (*Data groupings are labeled by letter on flow diagram*)

[B: Extract of patient record:
history,
past test results,
medications,
diagnosis],
[C: Discharge summary author ID, address],
[D: GP ID, address],
[E: Nominated 3rd party ID, address])

(**Data storage sites**:

[hospital storage:
paper patient record,
electronic patient record,
transcriptionist audio cassette library,
medical records department],
[GP office storage:
paper patient record,
electronic patient record])

Possible Transaction and Data Flows

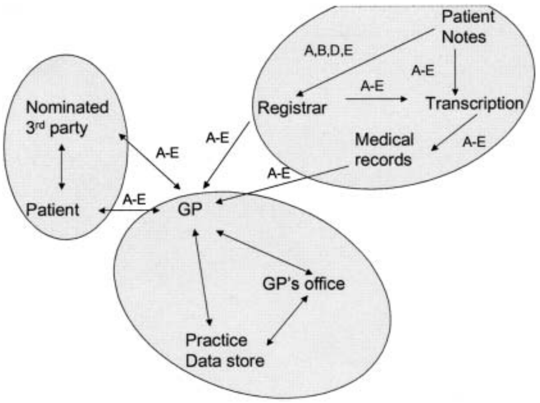

Any of the following combinations of communication channel, communication service, and communication device may be involved in transmission of some data for some of the transactions required during this scenario.

[Channel	**Service**	**Device]**
[Physical network:		
	Post	Letter
	Hospital internal courier	Audio cassette]
[Telephone network		
	Fax	Fax machine
	Internet e-mail	Computer
		Palm-top computer
	Internet IP, e.g., Web forms	Computer
		Palm-top computer]
[Internet/intranet		
	Internet e-mail	Computer
		Palm-top computer
	Internet IP, e.g., Web forms	Computer
		Palm-top computer]

Need-to-know Matrix (Agent vs. Data to determine need to access and use specified data):

	Patient ID	Author ID	Record Summary	GP ID	3rd Party ID
Patient	Y	Y	Y	Y	Y
Registrar	Y	Y	Y	Y	Y
Patient-specified party	?	?	?	?	?
Specialist	Y	Y	Y	Y	?
Medical transcriptionist	Y	Y	Y	Y	Y?
General practitioner (GP)	Y	Y	Y	Y	?
Practice nurse	Y	N	?	Y	?
Office clerk	Y	N	?	Y	?
Information system manager	Y	N	?	Y	?

Case 3—Hip Replacement

An elderly and independent woman who lives alone is admitted to the hospital for a hip replacement operation. Following her operation, she is transferred to a rehabilitation facility to improve her mobility before she returns home. When she is sufficiently mobile to return home, information is sent to her general practitioner, her local Aged Care Assessment Team, and home and community care services.

(**Agents**:

Patient,
[Patient nominated additional party *roles:*
patient relative/executor,
consulting clinician,
insurance company,
federal agency, e.g., immigration],
[Hospital *roles*:
clerical,
nursing team: NUM, assigned nurse,
surgical team: intern, resident, registrar, surgeon,
specialist clinicians: anaesthesia, physician, etc.],
[GP office *roles*:
general practitioner,
practice nurse,
office clerk,
information system manager],
[Rehab *roles*:
clerical,
nursing team: NUM, assigned nurse,
allied health: physiotherapy, occupational therapy, etc.,
clinical team: intern, resident, physician],
[Aged care assessment team])

(**Data**: [A: patient ID], (*Data groupings are labeled by letter on flow diagram*)

[B: Extract of patient record:
history,
past test results,
medications,
diagnosis,
surgical record],
[C: Discharge summary author ID, address],
[D: GP ID, address],
[E: Nominated 3rd party ID, address]
[F: Surgical Team ID])

(Data storage sites:

[hospital *storage*:
paper patient record,
electronic patient record,
transcriptionist audio cassette library,
medical records department],
[GP office *storage*:
paper patient record,
electronic patient record]
[rehab *storage*:
paper patient record,
electronic patient record,
transcriptionist audio cassette library,
medical records department],
[Aged care assessment *storage*:
paper patient record,
electronic patient record]),

Any of the following combinations of communication channel, communication service, and communication device may be involved in transmission of some data for some of the transactions required during this scenario.

[**Channel**	**Service**	**Device**]
[Face to Face:		
	Patient (as channel)	
	Referring Doctor (as channel)]	
[Physical network:		
	Post	Letter],
[Wireless local network		
	IR	Palm-top computer
		Printer
		Computer
		Courier bar code scanner]
[Wireless wide-area networks:		
	Paging	Pager
	Short message service	
	Mobile Phone	
	Voice	Mobile phone
	WAP	Mobile phone
		Palm-top computer]
[Telephone network		
	Voice	Telephone
	Fax	Fax machine
	Internet e-mail	Computer
		Palm-top computer
	Internet IP, e.g., Web forms	Computer
		Palm-top computer]
[Internet/Intranet		
	Test order entry/ reporting software	Computer
		Printer
		Storage archive, e.g., DAT
	Internet e-mail	Computer
		Palm-top computer
	Internet IP, e.g., Web forms	Computer
		Palm-top computer]