Toward a Standard for Community Participation in Terminology Development

Martin Benjamin[1]

[1]Kamusi Project International, Executive Director, Geneva, Switzerland

Abstract: This paper is intended to stimulate discussion about the creation of a standard for a community participation model of terminology development. It begins with a discussion of why the current state of terminology development in Africa brings forth the need for such a standard process. The paper then presents the outcome of an experiment in the development of ICT terminology for Swahili, as well as a subsequent project to incorporate the experiment's results in the creation of terminology sets in African languages for multiple domains. A final section sets forth considerations for developing a community standard, particularly for contexts where no corps of professional terminologists exists. The proposed process is expected to deliver linguistically-appropriate terminology that will be acceptable to specialists and the general public, and will be adopted within a wide range of technical activities.

Introduction: This paper is intended to stimulate discussion about the creation of a standard for a community participation model of terminology development. It begins with a discussion of why the current state of terminology development in Africa brings forth the need for such a standard process. The paper then presents the outcome of an experiment in the development of ICT terminology for Swahili, as well as a subsequent project to incorporate the experiment's results in the creation of terminology sets in African languages for multiple domains. A final section sets forth considerations for developing a community standard, particularly for contexts where no corps of professional terminologists exists. The proposed process is expected to deliver linguistically-appropriate terminology that will be acceptable to specialists and the general public, and will be adopted within a wide range of technical activities.

Terminology development in Africa: The emergence of a community participation model for terminology development is a response to the current haphazard state of the field in Africa. Approximately 2000 languages are spoken by about one billion people in Africa's 53[1] countries. Several dozen of these languages are each spoken by millions of people, often as official or national languages. Meanwhile, most technical activity is conducted in souvenir languages from the colonial era, particularly English, French, and Portuguese. Limited efforts for a few languages have been made to develop terminology that would be inclusive for the large majority of Africans who do not have advanced training in these foreign languages, notably for Swahili and for the official languages of South Africa. The state of terminology in South Africa, Tanzania, and Kenya has been the subject of research visits and interviews for the project under discussion. These projects have produced useful term sets, but also

---

[1] The number will have changed to 54 for those reading this paper after 9 July, 2011, the formal date of independence for Southern Sudan.

point to the weaknesses in current systems for the production and dissemination of terminology data. Because terminology development for most African languages is currently non-existent or in a nascent state, the moment is ripe to learn from past efforts and institute a standard process that can support a variety of languages and domains in diffuse political and infrastructure landscapes.

Terminology development in Africa is most systematic in South Africa. South Africa has 11 official languages, including English, Afrikaans, and nine languages that pre-date the colonial era. The Pan-South Africa Language Board (PanSALB) oversees language policy and coordinates development for the various languages, including convening National Language Boards for each official tongue. The NLBs operate independently when mandated to produce terminology sets, with the dual effect that terminologies for a domain might be produced for only some of the official languages, and that closely related languages such as Zulu and Xhosa may produce divergent terms for the same concept because the boards operate without reference to each other's activities. Terminology sets are deemed complete when they are approved by the language boards, with no opportunity for community review or comment. Terminology sets are generally presented as straight glosses of English source terms, without corresponding definitions of the concepts in the target language, thereby leaving potential users with few clues as to the meaning of terms that are being granted new senses or are being introduced to the language. Approved terminologies are published in print form only, and not widely available outside of certain libraries. Finally, though the government of South Africa owns the copyrights to the data, the prevailing national policy interprets that copyright in a restrictive manner that not only prohibits republication of terminology sets by others, but often inhibits its use within the relevant industries. In sum, South Africa has a refined infrastructure for expert development of terminology sets that produces official parallel glosses, without communication among terminologists working on related languages, without community validation and with a dissemination model at odds with maximization of community uptake.

The somewhat different situation for Swahili is complicated by a sense of ownership toward the language in two separate countries, Kenya and Tanzania, with no body that unifies language policy between them. Various agencies, committees, and academic departments take on the task of terminology development within different domains. This sometimes results in multiple glossaries for overlapping data; for example, BAKITA (the National Swahili Council of Tanzania), TUKI (Institute of Kiswahili Research at the University of Dar es Salaam), and Radio Tanzania developed three complete, and often conflicting, lists of country names. The most official process is for a term set to be approved by BAKITA. However, BAKITA is highly under-resourced, so is limited in the scope of work it can undertake. As with South Africa, terminology sets by BAKITA or others usually consist only of glosses, and are published only in print form without widespread distribution. Neither Tanzania nor Kenya impose the same copyright restrictions as South Africa, so opportunities for downstream use are in principle more open, when people are aware of the data's existence.

For other African languages and countries, the situation is generally much less formalized. Terminology work, when it occurs, can fall on academic units, or on interested parties such as health agencies. The results thus often do not receive the official imprimatur of a government body, but may become the de facto standard if uptake is sufficient and competing term sets are absent. Most projects

undertaken in this ad hoc fashion follow the model of expert production without a community validation procedure, straight term-for-term glosses without concept definitions, and limited dissemination.

Most prevalent, however, is a situation where terminology development does not occur in African languages. The primary reasons for this are:

1. Ambivalence by elites fluent in the souvenir languages toward the need for locally understandable terminologies
2. Paucity of resources to devote to linguistic issues in a tightly constrained budget environment
3. Lack of a process or mechanism to facilitate terminology development

This paper proposes the standardization of a process, along with an available terminology development mechanism. The acceptance of these features can reduce the costs of terminology development to a level manageable within many budgets, which in turn may create the conditions whereby domain-specific African language terminology becomes normalized as a goal of national policy planners.

The Maneno Magumu[2] experiment: In March 2010, an experiment in terminology development called "Maneno Magumu: Difficult Terms in Swahili ICT Localization" was conducted by the Kamusi Project, an NGO that produces digital resources for African languages including the "Internet Living Swahili Dictionary."[3] The purpose of the experiment was twofold. The immediate aim was to work toward a stable set of terms that can be used to localize Swahili software and ICT products, that will be acceptable to the Swahili linguistic community as they encounter expanded access to novel technologies. The larger goal was to gain experience bringing community members into the process of developing their own terminologies for domains of concern to them, with the intent of using the lessons of the experiment to construct a more permanent platform for terminology development for African (and potentially other) languages.

The Maneno Magumu experiment arose out of an ICT Terminology project undertaken for ANLoc, the African Network for Localization[4]. The ANLoc project involved the creation of terminology data sets for 12 African languages, centered around a core 2500 concepts that commonly appear in a wide range of software applications. ANLoc determined that standard ICT terminology was a necessary initial component of localization activities, to ensure that technology could be understood by its users in their own languages. Terms were selected based on a frequency analysis of their appearance in over 1100 software translation files, then given localization-specific definitions in English, as well as equivalent terms and definitions in French. African language partners then produced glosses and definitions for those terms in their languages, using on online interface called Glossmaster that was developed for the project by IT46[5]. This project was supported by IDRC, the International Development Research Center of Canada[6]. Groups were located around Africa that had the combination of ICT

---

[2] "Maneno Magumu" means "Difficult Words" in Swahili
[3] Internet Living Swahili Dictionary: http://kamusi.org; "kamusi" is the Swahili word for "dictionary," derived from Arabic "قاموس (kamus)," and cognate with many other African languages such as Hausa "k'aamùs."
[4] African Network for Localization: http://africanlocalization.net/
[5] Glossmaster: http://www.it46.se/glossmaster
[6] International Development Research Centre: http://www.idrc.ca

knowledge and linguistic competence for each language. Each group worked independently and produced a data set that has been made freely available for localization activities in those languages[7].

The ANLoc terminology project began by following a very traditional model of terminology development, what could be called the "experts-out" model. Experts-out terminology development begins with one or more specialists sitting in a room, looking at a vocabulary list for a particular domain, and saying, in effect, "For our language, the term for such-and-such a concept will henceforth be so-and-so." The reasoning behind the terms thereby devised may be impeccable, following all of the standard terminographic procedures - and the resulting dataset, as Microsoft learned, may still prove problematic[8]. The real test of a terminology set is not whether the terms are elegant or clever. The test is whether they are adopted by their intended use communities. As an example, the term "ngamizi" for "computer" was developed for the Microsoft Kiswahili localization glossary by a team of Swahili scholars, with a basis rooted in the poetic potentials of the language[9]; in the real world, the public ignored that coinage and instead use the term "kompyuta." As the ANLoc project moved along, and began reviewing terminology projects for other domains throughout Africa and elsewhere, it was increasingly understood that an experts-out approach often results in orphan terminologies that are not adopted by the public.

One of the languages in the ANLoc Terminology project was Swahili, the language at the core of the Kamusi Project. Terminology development was placed in the hands of the Tanzanian Linux Users Group (tzLUG), a group that consists mostly of ICT university students and recent graduates in Dar es Salaam. Part of the group's brief was to base their glossary as much as possible on the existing ICT terminologies developed for localization projects by Google (who were the first to introduce a widely-available Swahili interface[10]), Microsoft (who have now produced Swahili Language Interface Packs for three generations of Windows and Office[11]), and KiLinux[12] (a project based at the University of Dar es Salaam to localize Open Source software in Swahili, notably OpenOffice.org, an early incarnation of Firefox, and a never-completed attempt at the Linux operating system desktop), with the goal being a de facto standard terminology set[13] that could be accepted by all stakeholders. tzLUG produced rapid results that ended up being an inconsistent combination of terms from the various sources mixed with proposals from the young volunteers. Rather than release this data set as-is, thereby landing yet another competing set of terms into the Swahili localization environment, it was decided that the data

---

[7] ANLoc ICT Terminology Available to Download (Completed Terminology Sets): http://www.kamusi.org/en/ICT_terminology_downloads_10_languages

[8] "Why Microsoft Swahili Version Failed", Business Daily, 12 March 2008, archived at http://www.balancingactafrica.com/news/en/issue-no-396/computing/is-microsoft-swahili/en

[9] "'Ngamizi' was proposed by a Kenyan team member who referred to his first experience with a computer which processed (swallowed "-meza" in standard SW, "-miza" in dialect) data." Legère 2006, 182.

[10] Google in Swahili: http://www.google.com/webhp?hl=sw

[11] The most recent Microsoft Language Interface Pack, for Windows 7, is available at http://kamu.si/giGXKh

[12] The Kiswahili Linux Localisation Project, Department of Computer Science – UDSM: http://www.kilinux.udsm.ac.tz/

[13] Updegrove, Andrew, 1995. Consortia and the Role of Government in Standard Setting, in B. Kahin and J. Abbate (Eds.), *Standards Policy for Information Infrastructure* (pp. 321-350). Cambridge: MIT Press.

should be thoroughly reviewed in keeping with the aim of harmonizing ICT terminology for the language going forward.

Fortuitously, Microsoft approached the Kamusi Project in mid 2009 to review their existing 1500 term glossary and to participate in expanding it for 600 additional terms. Also in mid 2009, KiLinux released a long-awaited 1500 term glossary, expanded from the 700 term glossary they produced in 2005. In 2010, Google Nairobi agreed to share their internal glossary of more than 500 key terms. The Kamusi Project was thus able to compare all four data sets, and propose harmonized results that will not only be distributed by ANLoc, but are likely to be adopted in whole or in part by the prime movers in Swahili localization activities.

Some of the harmonization work was effortless, some was difficult, and more than a few terms initially proved impossible to reconcile across projects. When all the data sets were in agreement, such as the term "kiungo" for the noun concept "link," the term could be considered as established. In more difficult cases, research was able to reveal that one of the proposed Swahili terms had gained widespread usage while its competitors had failed to take hold, such as "Wavuti" having taken hold on the Internet as the term of choice for "Internet." Other conflicts could be resolved by reference to the guiding principle that the term should be as readily understandable as possible to the novice technology user, so that the word "umbizo" for the noun concept "format," derived from the Swahili root for "shape," was selected in favor of the Swahili-ized term "fomati" that has no resonance outside of the Swahili tech community. Some existing terms proved intractable, however. For example, "access" was variously rendered as "ufikivu," "mfiko," and "fikio" (all building from the same root, "kufika," meaning "to arrive"). In this example, no one term is obviously better than another, and corpus frequency analysis would be fruitless because some of the terms could appear in non-ICT contexts. Not being an official standard-making body[14], the Kamusi Project could not simply choose a term by fiat and mandate the other contenders to retranslate all their products. On the other hand, leaving all three terms in circulation can only lead to ongoing confusion by the user community; is the thing called "fikio" in Firefox the same or different from "mfiko" on Google websites or "ufikivu" in Microsoft Word? It is for these cases, as well as some of the new concepts identified by Microsoft, that the Maneno Magumu experiment was devised.

The experiment used a decidedly low-tech approach. With about 150 problematic terms, Maneno Magumu created six simple web forms with roughly 25 terms apiece. Each term was shown in English along with its English definition. Then each possible Swahili term that had been proposed was displayed, such as "ufikivu," "mfiko," and "fikio" from the example above. For each term, the participant had the option to propose their own Swahili equivalent if they were not satisfied with any of the existing options. Below the Swahili options were example sentences in English, if any existed in the source material. Finally, a text input box was provided for users to leave open-ended comments about the

---

[14] Hawkins, Richard, 2009. Business Models and the Dynamics of Supply and Demand for Standards, in K. Jakobs (Ed.), *Information Communication Technology Standardization for E-Business Sectors; Integrating Supply and Demand Factors* (pp. 31-51). Hershey: Information Science Reference.

term. In this way, participants were able to go through the experiment term by term, cast non-binding votes on their preferences, suggest new options, and/or offer their opinions and insights.

Invitations were circulated widely within the Swahili online community. Announcements were posted in Swahili and English to a comprehensive assortment of popular and academic Swahili mailing lists, the Swahili Wikipedia users group, Swahili groups on Facebook, blogs, Twitter, and the Kamusi Project homepage. It was stressed that the effort was purely voluntary (no pay or prizes) and that the initiative had a tight time limit. The time limit was deemed important in order to prevent the experiment from dragging indefinitely. The original seven days were stretched to fourteen, with follow-up notices to the various forums.

From these outreach efforts, about twenty members of the community took the time to complete one or more of the term packs, with at least ten users completing each pack. As responses arrived, patterns began to emerge:

- terms for which almost everyone agreed on a single answer
- terms for which community sentiment was divided among two or more acceptable answers
- terms for which the community did not like any of the suggested answers. This category was further divided into:
    - terms for which community members offered better answers
    - terms for which the community remained stumped

When participants trended toward a single answer, the conclusion was to designate that as the term to use going forward, without regard for the provenance of various terms. When the community was divided among terms that were deemed generally acceptable, further analysis was necessary in order to nominate a term as the preferred Swahili equivalent. Four primary criteria were used:

- Frequency analysis of the competing terms within Google search results. If search strings were not inherently unique to Swahili, they were appended with "kwenye," a word meaning with/ at/ in/ to, that is essentially exclusive to Swahili and appears in almost every Swahili document. Some terms showed one candidate with thousands or hundreds of thousands of results, with the competing term yielding almost none. In some cases, however, the terms had similarly large numbers of results, or similarly few.
- Existing usage within localization. All other things being equal, preference was given to terms that were already in use within major software products. Microsoft terms were given particular primacy because that company's software is the most difficult to change; their products are only re-translated once every three or four years, and users are not likely to update their versions even that often, so terms within the Microsoft system tend to remain there for a very long time. (However, no Microsoft term was kept that was shown to be less preferable due to other criteria; about 1/3 of their existing terms were recommended to be changed.) Google terms were deemed easier to change than Microsoft terms, since the interface is entirely online and often can be converted instantaneously. Google proposals had a generally high basis in combining the technical and linguistic concepts underlying a term, but fewer of them were

available for consideration, and their list was far from infallible. (Google terms regarding Search were given highest regard, since their interface controls about 100% of the Swahili market.) KiLinux terms were preserved when they accorded with popular sentiment or other criteria, or were the browser-centric terms chosen for Firefox, the only browser localized to Swahili at that time.

- Uniqueness and/or ease of use within a localization context. For example, "taarifa" was a more popular choice for "notification" than was "arifa," but "arifa" was selected for this specific concept because "taarifa" is a broad term for "information" that appears in numerous other uses. Similarly, "upeo" was chosen for "maximum" in favor of longer, more popular options ("kiwango cha juu" and "isiyozidi") because it is the simplest to use grammatically, and also because, with only four letters, it will fit most easily as a replacement for the three-letter abbreviation "max" in situations where space is limited. Although community participation provided an insight into popular opinion, the experts' actual experience with Swahili localization issues was sometimes given extra weight in the final decision.

- Swahili origins. Generally, a term that can be traced to a root concept in Swahili was given preference to a Swahili-ized version of an English term. For example, the Swahili word "tarakimu" translates as "numeral" or "digit," so was deemed a more understandable basis for "digital" than any of the attempts in the existing Microsoft glossary (dijito, digito, and dijiti). However, Swahili origins were not given precedence if other criteria pointed toward an English-derived term. For example, "faili" was selected for "file" as the clear community favorite, the most frequent research result, and the more well-established term within localization, even though the competing term "jalada" has a better Swahili pedigree.

The community input process was especially helpful in some cases where none of the existing proposed terms were found acceptable. Community members provided new analysis about the deeper meanings of problematic terms, leading to new glosses that would never have arisen without their insights. An impasse about the term "clipboard," for example, was resolved when a participant commented that the item's function is actually to copy something to a figurative board, rather than to clip it; his proposed term, "ubao nakili," is an elegant solution that translates back to English as "copying board." Similarly, a participant suggestion of "hifadhi muda" for "cache" perfectly encapsulates the concept of a cache as temporary storage.

In a few cases, even the community input process failed to yield any useful suggestions. These problem terms were taken to private discussions among "the experts," individuals who have worked extensively on Swahili technical translations. A 40 minute discussion via instant message with Google's Swahili Language Specialist in Nairobi, for example, tore apart the inner meaning and debated possible glosses for "mode," until the term "hali-tumizi" (more or less "usable condition") jumped out. Other terms developed in this way include "kompyuta ya kibao" for "tablet computer," and "kiguso" for "touchpad." These cases show that democratic participation cannot provide final answers for all development of technical terminology. However, exclusive recourse to the experts was a last resort that only proved necessary in a few extraordinary cases.

The Maneno Magumu experiment demonstrated that the collective mind offers seven important advantages to the process of terminology development:

1. Community input can validate proposed terms by confirming that they are understandable in the domain context.
2. Community response can indicate when terms are unacceptable or cannot be understood.
3. When several acceptable terms are under consideration, the weight of community responses can show which term is the most preferred.
4. Participants can propose terms that are better than those presented by the experts.
5. Participants can suggest approaches for thinking about a concept that stimulate the experts or other community members to devise an appropriate term.
6. Sometimes community feedback yields no actionable information, which highlights the most intractable terms on which the experts need to concentrate their efforts.
7. By inviting community members behind the curtains of terminology development, the process becomes much more transparent, therefore increasing the likely acceptance, ownership, and use of the resulting data.

Rather than "experts-out" terminology development, the model thus developed is "experts-with" - a system in which experts employ their usual tricks of language analysis, corpus-based research, and poetic coinage, and then admit that they do not have all the answers and open the door to community participation when they run into difficulties. In most cases, community input, often as a response to proposals from subject specialists, provided crucial guidance toward evaluating what will be understandable to the millions of people who will eventually make use of information technology localized into Swahili. In this way, the Maneno Magumu experiment demonstrates the efficacy of incorporating pseudo-democratic community participation within the process of terminology development.

KamusiTERMS: The positive outcome of the Maneno Magumu experiment has led to the creation of KamusiTERMS, the Kamusi for Technology, Economy, Rights, Medicine, and Science. KamusiTERMS is intended to become a multilingual, multi-domain development platform and term bank for African languages. It will be built on the proposed experts-with process of terminology development. The design of the KamusiTERMS software pays special attention to community participation features. Experts will easily flag entries for public review. Community members will be able to offer comments on proposals, suggest their own glosses or definitions, and cast non-binding votes on competing proposals. Project experts will then review the community feedback and make final decisions on difficult terms.

Terminology sets that are validated through this community participation process will be made freely available to the public through several electronic mechanisms. Data sets will be fully searchable on the web, using mobile technologies such as SMS and smartphones, and with special offline platform-independent software. Users will also be able to download terminology sets as text or PDF files that can be printed for use when electronic devices are not available.

It is hoped that KamusiTERMS will become a central terminology repository for any number of domains, for numerous languages of Africa. Development of the software for this initiative is expected to be complete in late 2011, after which it is anticipated that testing and refinement will occur with multilingual terminology projects for ICT and health domains. Official language bodies will be invited to participate in the development or review of terminology sets, with the goal of arriving at a system that produces standard terms with both community validation and approval by official language standards bodies.

Toward a proposed standard for participatory community terminology development: The community participation model is especially important for countries and languages that do not have a cadre of professional terminologists, a condition that describes most of Africa. Even languages overseen by bodies such as PanSALB and BAKITA stand to benefit from joining such a standard, adherence to which will facilitate terminology development through wider input and validation, and consequently lead to greater uptake.

The purposes of including linguistic user communities within the process of developing terminologies are:

1. To validate terms that are proposed by experts
2. To indicate preferences among multiple proposed terms
3. To generate suggested terms or approaches in cases where experts have no proposals
4. To augment community buy-in to the final terminology set

Therefore, a number of practices should become a standard part of the terminology development process:

1. Expert groups that consist of both language and subject specialists should collaborate in the development of proposed terms for a terminology set, using accepted best practices for terminology development[15].
2. All terms should receive both a gloss and a definition in the target language by the end of the process.
3. If experts have any doubt about a proposed term, they should make it open to community review.
4. If experts identify two or more potential glosses for a term, those terms should be put to the community for a non-binding vote.
5. Terms for which experts cannot generate satisfactory proposals should be put to the community for discussion.
6. "The community" should be specifically constituted for each terminology set by actively inviting stakeholders for the relevant domain, including traditional and social networks, and training them in the methods of participation.

---

[15] UNESCO / Prepared by Infoterm, 2005, Guidelines for Terminology Policies: Formulating and implementing terminology policy in language communities. UNESCO: Paris.

7. Expert groups should not be bound by numerical adherence to the results of community voting, but should use the community input as guidance to inform their subsequent debates and decisions.
8. The decision-making process should be made transparent to the community, through available records about why the expert groups selected one term over another.
9. Resulting terminology sets should be made available to the community on a free and open basis, including electronic distribution in an accessible location and a copyright that maximizes sharing and reusability.

Community participation therefore means democratic input, but not majority rule. The process recognizes that experts tend to know more about a subject, or about linguistic subtleties, than the public at large. The process also acknowledges that decisions will need to be made that do not satisfy all parties, and that an expert group acting with advice from the community is best positioned to make those decisions. However, the process also admits that experts do not have all the answers, and accepts that an appeal to the collective mind can often generate better results than an exercise conducted by a few individuals behind closed doors.

Conclusions: Terminology sets that result from the community participation process will benefit from the input and validation of their intended users. Participatory review will increase the likelihood that the terminology will be widely accepted and used. Because terminology is the interface for communication within a domain, community investment in and ownership of their terms of discourse will enable technical activity within languages that successfully engage in its production.

This paper does not propose a final standard for community-enhanced terminology development. Rather, it has raised several issues of concern, with particular reference to the African context. Through an experiment with participatory approaches in the production of ICT terminology for Swahili, and the development of a mechanism for wider implementation, it has been possible to elaborate many of the considerations that should go into the creation of a standard[16]. In keeping with the participatory ethos that underlies the proposed process, it is hoped that readers will offer feedback that can be cycled into a consensus that can become the standard for community participation in terminology development.

References:

Folmer, Erwin, and Jack Verhoosel, 2011. State of the Art on Semantic IS Standardization, Interoperability and Quality, University of Twente.

Hawkins, Richard, 2009. Business Models and the Dynamics of Supply and Demand for Standards, in K. Jakobs (Ed.), *Information Communication Technology Standardization for E-Business Sectors; Integrating Supply and Demand Factors* (pp. 31-51). Hershey: Information Science Reference.

---

[16] Folmer, Erwin, and Jack Verhoosel, 2011. State of the Art on Semantic IS Standardization, Interoperability and Quality, University of Twente.

Karsten Legère (2006). Formal and Informal Development of the Swahili Language: Focus on Tanzania, Selected Proceedings of the 36th Annual Conference on African Linguistics: Shifting the Center of Africanism in Language Politics and Economic Globalization. Olaoba F. Arasanyin and Michael A. Pemberton (eds.). Somerville, MA: Cascadilla Proceedings Project, pp. 176-184.

UNESCO / Prepared by Infoterm, 2005, Guidelines for Terminology Policies: Formulating and implementing terminology policy in language communities. UNESCO: Paris.
Updegrove, Andrew, 1995. Consortia and the Role of Government in Standard Setting, in B. Kahin and J. Abbate (Eds.), *Standards Policy for Information Infrastructure* (pp. 321-350). Cambridge: MIT Press.