NORTH CAROLINA JOURNAL of
LAW & TECHNOLOGY
Volume 28, Issue 1 2026

ARTICLE

FORGIVE ME, ALGORITHM: MAGISTERIUM AI, THE CLERGY PENITENT PRIVILEGE, AND THE CASE FOR A STATUTORY RELIGIOUS AI PRIVILEGE

Casey Rockwell† & Michael Conklin‡

"It's never going to give you a wrong or false answer . . . ."
— Father Philip Larrey, chair of the Product Advisory Board of Magisterium AI[1]

The intersection of artificial intelligence and religious practice has produced an unexamined gap in U.S. evidentiary law. Magisterium AI is a chatbot trained exclusively on official Catholic Church documents and it is now integrated into consumer applications in more than 165 countries. It represents a new category of technology: the institutional religious AI chatbot. When users disclose personal moral and spiritual information to such a platform, no U.S. evidentiary privilege protects those communications. Instead, the clergy-penitent privilege, codified in all fifty states and grounded

© 2026 Casey Rockwell & Michael Conklin.

† JD, EdD; Assistant Director of the School of Business, University of Arkansas at Little Rock.

‡ JD, LLM, MBA, MSBA, MPhil; Assistant Professor of Business Law, Texas A&M University Central Texas; Lecturer, Texas A&M University School of Law.

1. Claire Giangravé, *Artificial Intelligence Program Intelligence Program Poised to Shake up Long-Held Catholic Doctrine*, AM. JESUIT REV. (Aug. 25, 2023), https://www.americamagazine.org/faith/2023/08/25/artificial-intelligence-vatican-theology-245929/ [https://perma.cc/38UT-S9GR].

in the Free Exercise Clause, requires a human spiritual advisor as a party to the communication; the Third-Party Doctrine makes AI-mediated communications presumptively accessible in the discovery process; and the recent federal decision in United States v. Heppner *forecloses privilege claims over commercial AI platform communications. This Article argues that the resulting legal gap is both doctrinally anomalous and practically harmful. Drawing on the privilege-recognition framework of* Jaffee v. Redmond *and Federal Rule of Evidence 501, the canonical* sigillum sacramentale *under the 1983 Code of Canon Law, the Vatican's 2025 document Antiqua et Nova, and state mandatory-reporting legislation placing pressure on traditional confessional practice, this Article proposes a novel, "Digital Seal" statutory framework. The framework would extend a qualified evidentiary privilege to any AI platform that meets the following three conditions: institutional anchoring within a recognized religious organization, a defined data-minimization architecture, and canonical or equivalent ecclesiastical oversight. It is designed to be confessionally neutral and applicable across religious traditions.*

Table of Contents

I. Introduction 87
II. Magisterium AI and the Institutional Religious AI Landscape 90
A. Architecture and Design 90
B. The Broader Phenomenon: Religious AI Beyond Catholicism 91
III. The Clergy-Penitent Privilege in U.S. Law 92
A. Historical and Constitutional Foundations 92
B. Statutory Formulations and Common Elements 94
C. The Canonical Dimension: The Sigillum Sacramentale 94
D. The Privilege Under Legislative Pressure: Washington SB 5375 95
E. Applicability to Mistaken Identity 97
IV. Three Independent Doctrinal Barriers to Privilege Recognition 99
A. The Human-Advisor Requirement 99
B. The Third-Party Doctrine 101
C. United States v. Heppner and the Foreclosure of AI Privilege Claims 105
V. The Scholarly Debate over an AI Interaction Privilege 106

VI. THE DIGITAL SEAL: A PROPOSED STATUTORY FRAMEWORK . 110
A. *Design Principles* 110
B. *The Three Structural Conditions* 111
C. *Scope, Exceptions, and the Mandatory Reporting Problem* 114
D. *Canonical Alignment, Constitutional Considerations, and the Limits of Civil Law* 116
E. *Comparative Law Grounding and Model Statutory Text* 118
VII. CONCLUSION 121

I. INTRODUCTION

U.S. evidentiary law protects certain relationships from compelled disclosure in judicial proceedings. Attorney-client, physician-patient, therapist-patient, and clergy-penitent communications all receive this protection.[2] In each case the rationale is the same: the social benefit of candid communication within that relationship outweighs the evidentiary cost of keeping it out of court.[3]

Artificial Intelligence (hereinafter "AI") has created a problem that existing privilege law is ill-equipped to handle. Users now share with AI platforms the same kind of sensitive information they would share with an attorney, a doctor, a therapist, and a priest. If that information were disclosed to an actual human professional, privilege would protect it. Disclosed to religious AI instead, privilege would not.[4]

---

2. *See* Jaffee v. Redmond, 518 U.S. 1, 9–10 (1996) (articulating that evidentiary privileges protect relationships whose social value depends on candid, confidential communication).

3. *Id.* at 11; *see also* Diane Marie Amann & Edward J. Imwinkelried, *The Supreme Court's Decision to Recognize a Psychotherapist Privilege in* Jaffee v. Redmond, *116 S. Ct. 1923 (1996): The Meaning of the Term 'Experience' and the Role of 'Reason' Under Federal Rule of Evidence 501*, 65 U. CIN. L. REV. 1019, 1021 (1997) (explaining that privileges obstruct the search for truth and therefore the Supreme Court has historically recognized a new privilege only when the public good it serves is great enough to override the interest in ascertaining truth, and noting that *Jaffee* nonetheless recognized a psychotherapist-patient privilege).

4. Samuel N. Dick, *Virtual Confessions: Examining the Clergy Privilege's Extension to Artificially Intelligent Religious Robots*, 9 U. CIN. INTELL. PROP. & COMPUT. L.J. 76, 78–80 (2024).

This Article examines that problem in one specific context: the institutional religious AI chatbot. The clearest example is Magisterium AI, created by the technology company Longbeard and trained on official Catholic Church documents including the Code of Canon Law, papal encyclicals, the Catechism of the Catholic Church, and conciliar decrees.[5] The platform now operates in more than 165 countries and is integrated into the Catholic prayer app Hallow, where it reaches a large lay audience.[6] Similar platforms are being developed for other faith communities.[7]

The question this Article poses is narrow: when a Magisterium AI user discloses personal moral or spiritual information, does the clergy-penitent privilege (i.e., the priest-penitent privilege) protect that communication? Under current doctrine, the answer is no. This Article argues that it should under defined conditions, and proposes a "Digital Seal" statute to achieve that result.

Two clarifications are necessary. First, this Article takes no position on whether AI-mediated spiritual guidance is theologically appropriate. That question belongs to religious institutions, not civil law.[8] The Vatican's 2025 document Antiqua et Nova addresses that institutional question; this Article draws on it only for its structural

5. Kenny Ang, *Magisterium AI in Theological Inquiry and Religious Education: Challenges and Emerging Horizons*, 13 SCIENTIA ET FIDES 239, 240–41 (2025). The platform's training corpus is described as encompassing documents within both the Ordinary and Extraordinary Magisterium, including Scripture, the Catechism, papal encyclicals, conciliar decrees, and the writings of recognized Doctors of the Church.

6. *Magisterium AI FAQ*, HALLOW HELP & SUPPORT CTR. (Mar. 24, 2026), https://help.hallow.com/en/articles/10601094-magisterium-ai-faq [https://perma.cc/43QA-YUAQ].

7. Dick, *supra* note 4, at 77–78 (surveying AI platforms serving Jewish, Islamic, Protestant, and Buddhist communities in roles including sermon preparation, spiritual counseling, and sacramental preparation).

8. *See* Dicastery for the Doctrine of the Faith & Dicastery for Culture and Education, *Antiqua et Nova: Note on the Relationship Between Artificial Intelligence and Human Intelligence* ¶ 112 (Jan. 28, 2025) [hereinafter *Antiqua et Nova*], https://www.vatican.va/roman_curia/congregations/cfaith/documents/rc_ddf_doc_20250128_antiqua-et-nova_en.html [https://perma.cc/L3KL-DVRL] (reserving theological and moral judgment on AI's role in religious life to institutional discernment guided by human authority rather than civil law).

implications for ecclesiastical accountability.[9] Second, this Article does not claim Magisterium AI functions as a spiritual counselor in any sacramental sense. Here, the claim is simpler: users do in fact share personal moral and spiritual information with these platforms, the documentation of that behavior is substantial, and the law currently offers no protection for those disclosures.[10]

The canonical foundation for the Digital Seal framework was further strengthened in 2026 when Pope Leo XIV issued the first papal social encyclical devoted directly to artificial intelligence, *Magnifica Humanitas*.[11] The encyclical emphatically warned about the potential distortive and dehumanizing effects of AI, and called for regulatory tools to promote justice.[12] These principles directly support two core structural components of the Digital Seal framework proposed in this Article: institutional ecclesiastical oversight of religious AI systems and data-minimization requirements designed to prevent spiritual disclosures from being transformed into commercial analytics or training data.

Part II describes Magisterium AI and the wider religious AI landscape. Part III examines the clergy-penitent privilege, its statutory forms, and the canonical law of the *sigillum sacramentale*. Part IV identifies the three doctrinal barriers that block privilege protection for AI-mediated religious communications: the human-advisor requirement, the Third-Party Doctrine, and *United States v. Heppner*. Part V surveys the academic debate over an AI interaction privilege. Part VI presents the Digital Seal framework. Finally, Part VII

9. *Antiqua et Nova*, *supra* note 8, at ¶¶ 42, 112.

10. Robert Messer & Kathrin Zeilmann, *Chatting to Jesus on Your Phone: What's the Church's Stance on AI*, Deutsche Presse-Agentur (May 11, 2026, at 08:08 ET), https://www.yahoo.com/news/articles/chatting-jesus-phone-whats-churchs-120851760.html [https://perma.cc/7P32-PNAK]. This source is cited solely as a factual predicate documenting widespread consumer deployment of institutional religious AI platforms and the questions such deployment raises; no analytical proposition in this Article rests upon it.

11. Pope Leo XIV, *Magnifica Humanitas*: On Safeguarding the Human Person in the Time of Artificial Intelligence (May 15, 2026) (encyclical letter), https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html [https://perma.cc/4V3N-V6FP].

12. *Id.* at ¶¶ 5, 10, 15.

concludes by summarizing the case for a Digital Seal exception and emphasizing the rapidly evolving nature of the issue.

II. MAGISTERIUM AI AND THE INSTITUTIONAL RELIGIOUS AI LANDSCAPE

A. Architecture and Design

Magisterium AI launched in July 2023.[13] Its founder describes the platform's goal as making "Catholic insight and Church teaching available to anyone in the world, on any device, in their native language."[14] Unlike general-purpose AI systems trained on broad internet data, Magisterium AI draws only from a curated set of official Church documents.[15] A scholarly advisory committee chaired by the Rector of the Pontifical Oriental Institute in Rome oversees that corpus.[16] The company's "Why Magisterium" page describes its functionality as, ". . . fundamentally different."[17] It learns exclusively from a meticulously curated database of authoritative Catholic sources–over 23,000 official Magisterial documents (like the Catechism, Papal Encyclicals, Council Decrees, Canon Law) and over 2,300 foundational Scholarly works (from Church Fathers, Doctors like Aquinas and Augustine, the Bible, and key commentaries)."[18] Magisterium defines uses as, "Getting accurate answers to your specific questions about faith and morals . . . [and] clarifying complex doctrines (like the Trinity or Sacraments) faithfully."[19]

---

13. Ang, *supra* note 5, at 241.

14. *Id.* (quoting Matthew Sanders, founder and CEO of Longbeard, on the platform's mission statement).

15. Giangravé, *supra* note 1("Unlike other AI programs, which have access to vast swaths of ever-evolving data, the information used by Magisterium AI is limited to official church documents and is carefully curated.").

16. *Id.* (describing the scholarly advisory committee chaired by the Rector of the Pontifical Oriental Institute in Rome and its role in curating the training corpus).

17. *Why Magisterium AI?*, MAGISTERIUM AI, https://www.magisterium.com/about/why-magisterium-ai [https://perma.cc/5EFN-328P] (last visited Aug. 28, 2026).

18. *Id.*

19. *Id.*

The platform's published FAQ states that it "is not meant for spiritual or moral direction" and directs users with personal or pastoral questions to consult a priest.[20] That disclaimer matters legally, but it does not resolve the privilege question. The disclaimer speaks to intended use. Actual use, as the literature on AI-mediated communications shows, regularly involves personal moral and spiritual disclosure that users treat as confidential.[21]

B. The Broader Phenomenon: Religious AI Beyond Catholicism

Magisterium AI is the best-documented example of a wider pattern. AI platforms serving Jewish, Islamic, Protestant, and Buddhist communities are either live or in development, each raising the same privacy questions.[22] Within the Jewish faith, Rebbe.IO was established in 2023 to allow users to "[a]sk anything about Torah, Halacha, the holidays, and Jewish life. Get clear, instant answers grounded in classic sources—always free."[23] Within the Protestant-evangelical faith, BibleChat claims market share of over 10 million users [24] offering, "the world's most popular faith app, helping Christians find peace and support every day through God's Word."[25] The legal analysis in this Article is not specific to Catholicism. The Digital Seal framework proposed as a solution in Part VI applies to any religious tradition whose AI platform meets the structural requirements.

This Article relies heavily on Catholic sources because Magisterium AI is the most developed and documented platform available. The Catholic canonical tradition, particularly the sigillum sacramentale, also provides the most detailed institutional model for

20. *Magisterium AI FAQ*, *supra* note 6.

21. Dick, *supra* note 4, at 95–96 (documenting patterns of personal moral and spiritual disclosure to AI platforms across quasi-pastoral use cases).

22. *Id.* at 77–78.

23. *Rebbe.IO — World's First AI Rabbi*, REBBE.IO, https://rebbe.io/ [https://perma.cc/LG24-EFWA] (last visited Aug. 28, 2026).

24. *Romanian Startup Bible Chat Secures $14M Series A Led by True Ventures*, VESTBEE (Feb. 20, 2025), https://www.vestbee.com/insights/articles/bible-chat-secures-14-m [https://perma.cc/PQE4-XVE3].

25. *About Bible Chat — Explore Our Mission, Team, and Story*, BIBLE CHAT, https://thebiblechat.com/about-us/ [https://perma.cc/RQM9-TK8A] (last visited Aug. 28, 2026).

confidentiality obligations in religious communication. The same underlying legal questions apply equally to AI platforms serving other faiths.

III. The Clergy-Penitent Privilege in U.S. Law

A. Historical and Constitutional Foundations

The clergy-penitent privilege is one of the oldest evidentiary privileges in Anglo-American law.[26] Its U.S. origins trace to the 1813 case of *People v. Phillips*.[27] There, Phillips and his wife were charged with receiving stolen goods, but claimed they had received restitution from their pastor.[28] The pastor was then asked to disclose what information he had.[29] Citing specific provisions of the Catholic confession, the court recognized the sacrament and secrecy required and, should the priest be required to disclose, the priest would be breaking the seal.[30] The court further reasoned that a balance must be struck between the New York Constitution and the First Amendment protection of religious freedom, indicating that the two cannot coexist when the rule to discuss punishes the practice of one's religion.[31] Therefore, the court refused to compel a Catholic priest to disclose what he had

26. Dick, *supra* note 4, at 79–81 (tracing the historical development of the clergy-penitent privilege in English and American evidentiary law).

27. People v. Philips, (N.Y. Ct. Gen. Sess. 1813), *discussed in* Ann M. Murphy, *Washington State Settles Controversy over Child Abuse Law That Tested the Limits of Priest-Penitent Privilege*, Religion Unplugged (Oct. 30, 2025), https://religionunplugged.com/news/2025/10/30/washington-state-settles-controversy-over-child-abuse-law-that-tested-the-limits-of-priest-penitent-privilege [https://perma.cc/V786-UA8J]. People v. Philips is not officially reported, but was set forth by the attorney who participated in the case as amicus curiae, and later reprinted in I Western L.J. 109 (1843), among other reprintings. The original records of the case are on display in the library of the Court of General Sessions of the County of New York.

28. People v. Philips (N.Y. Ct. Gen. Sess. 1813) (DeWitt Clinton, Mayor), *reprinted in* William Sampson, The Catholic Question in America, 4–5 (De Capo Press 1813).

29. *Id.* 4, 8.

30. People v. Philips (N.Y. Ct. Gen. Sess. 1813), *reprinted in Privileged Communications to Clergymen*, 1 Cath. Law. 199, 202 (1955).

31. *Id.* at 206–207.

learned in confession.[32] Today the privilege is recognized in all fifty states, though it varies in scope by jurisdiction.[33]

The First Amendment's Free Exercise Clause provides the constitutional basis for the clergy-penitent privilege.[34] Courts have consistently held that the Clause protects religious institutions' internal governance from civil-law intrusion.[35] Forcing disclosure of sacramental communications would chill religious practice and directly interfere with institutional religious affairs.[36]

The clergy-penitent privilege also has a common-law basis independent of the Constitution. In *Jaffee v. Redmond*, the Supreme Court explained that all privileges protect relationships producing "public good" that depends on candid communication.[37] It rationalized that the treatment depends on the transparency that would not occur without confidentiality. The Court further refused to make this application on a case-by-case basis, as their determined that the lack of a consistent standard offers little security to the person confiding.[38]

---

32. People v. Philips (N.Y. Ct. Gen. Sess. 1813), *discussed in* Ann M. Murphy, *Washington State Settles Controversy over Child Abuse Law That Tested the Limits of Priest-Penitent Privilege*, RELIGION UNPLUGGED (Oct. 30, 2025).

33. *See id.*; Dick, *supra* note 4, at 81 (noting that the clergy-penitent privilege "is a right recognized in all 50 states").

34. U.S. Const. amend. I.

35. *See* Our Lady of Guadalupe Sch. v. Morrissey-Berru, 591 U.S. 732, 746 (2020) (reaffirming the ministerial exception doctrine, which shields religious institutions' employment decisions regarding certain employees from civil-law scrutiny); *see also* Richard C. Schragger, Micah Schwartzman & Nelson Tebbe, *Reestablishing Religion*, 92 U. CHI. L. REV. 199, 208–09 (2025) (surveying the Supreme Court's recent doctrinal structure for religious institutional autonomy).

36. Mockaitis v. Harcleroad, 104 F.3d 1522, 1530 (9th Cir. 1997) (holding that Oregon district attorney's secret recording of a jailhouse confession between a priest and an inmate violated the First Amendment's Free Exercise Clause).

37. Jaffee v. Redmond, 518 U.S. 1, 9–10 (1996); *see also* Robert H. Aronson, *The Mental Health Provider Privilege in the Wake of* Jaffee v. Redmond, 54 OKLA. L. REV. 591, 593–94 (2001) (noting that almost every state recognizes some form of religious privilege, that the communicant ordinarily holds it, and that the Uniform Rules of Evidence kept it unchanged while declining to add new privileges because privileges obstruct the search for probative evidence).

38. *Jaffee*, 518 U.S. at 17–18.

The clergy-penitent relationship clearly fits this description,[39] and, therefore, it is precisely what religious AI platforms unsettle as users bring the same reliance on confidentiality to a chatbot back by privacy policy rather than privilege.

B. Statutory Formulations and Common Elements

State statutes share three basic elements, though the specific language varies. First, the communication must be made to a person recognized as clergy or a minister of religion. New York provides that "a clergyman, or other minister of any religion or duly accredited Christian Science practitioner, shall not be allowed to disclose a confession or confidence made to him in his professional character as spiritual advisor."[40] California defines the privilege-holder as "a priest, minister, religious practitioner, or similar functionary of a church or of a religious denomination or religious organization."[41]

Second, the communication must be made in the clergy member's professional capacity as a spiritual advisor, not in some other role.[42]

Third, the communication must be intended as confidential.[43] States differ on who holds the privilege: some place it solely with the clergy member, others with the penitent, and others with both.[44] This jurisdictional split is not discussed further in this Article because whether any privilege attaches at all to AI-mediated communications is resolved at a prior step by the human-advisor requirement.

C. The Canonical Dimension: The Sigillum Sacramentale

Catholic canon law imposes an obligation that predates and parallels the civil-law privilege. Canon 983 § 1 of the 1983 Code of Canon Law states: "The sacramental seal is inviolable; therefore it is absolutely forbidden for a confessor to betray in any way a penitent in

39. Dick, *supra* note 4, at 80–82.

40. N.Y. C.P.L.R. § 4505 (McKinney 2025).

41. CAL. EVID. CODE § 1030 (West 2026).

42. Dick, *supra* note 4, at 82–83.

43. *Id.*

44. *Id.* at 84 (surveying state statutory variations on privilege-holder identity and the controversy over whether the privilege belongs to the penitent, the clergyman, or both).

words or in any other manner and for any reason."[45] Canon 1386 §1 subjects any priest who violates the seal directly to automatic excommunication, a *latae sententiae* penalty requiring no trial.[46]

In 2019, the Apostolic Penitentiary reaffirmed the seal's absolute character in its Note on the Importance of the Internal Forum and the Inviolability of the Sacramental Seal, which was issued in response to legislative efforts in several jurisdictions to compel confessional disclosures.[47] The Note stated that the seal "admits of no exceptions" and "cannot be subject to any pressure, including the pressure of civil law."[48]

The canonical structure of the *sigillum* is relevant here for two reasons. First, it demonstrates that religious institutions have already developed sophisticated, legally cognizable frameworks for confidential spiritual communications, rooted in institutional accountability and canonical sanction. Second, it shows that the confidentiality interests at stake in religious AI communications are not novel claims. They reflect centuries of institutional practice that civil law has consistently respected. The Digital Seal framework extends that respect to a new technological context.

D. The Privilege Under Legislative Pressure: Washington SB 5375

Mandatory child abuse reporting requirements have put increasing pressure on the traditional privilege. In May 2025, Washington enacted Senate Bill 5375, designating clergy as mandatory reporters of child sexual abuse, physical abuse, and neglect, including

---

45. *Codex Iuris Canonici* 1983 CODE c.983, § 1.

46. *Codex Iuris Canonici* 1983 CODE c.1386 § 1. A *latae sententiae* penalty operates automatically upon commission of the prohibited act, without any judicial process, declaration, or sentence by an ecclesiastical authority.

47. Apostolic Penitentiary, *Note of the Apostolic Penitentiary on the Importance of the Internal Forum and the Inviolability of the Sacramental Seal*, HOLY SEE (June 29, 2019) [hereinafter *Apostolic Penitentiary Note*], https://www.vatican.va/roman_curia/tribunals/apost_penit/documents/rc_trib_appen_pro_20190629_forointerno_en.html [https://perma.cc/LQE8-XRUQ].

48. *Id.* (stating that the seal of confession "admits of no exceptions" and "cannot be subject to any pressure, including the pressure of civil law," because it "protects the freedom of the person who approaches the Sacrament of Penance").

disclosures made within a confidential sacred rite.[49] The statute triggered immediate constitutional controversy under the framework of *Employment Division v. Smith* and *Fulton v. City of Philadelphia*.[50] *Smith* holds that neutral laws which incidentally burden religion need to give way. *Fulton* holds that laws which are not neutral, including but not limited to cases that carve out secular exceptions while denying comparable religious one, face strict scrutiny.[51]

Most states have handled this tension by preserving the clergy privilege while requiring reporting of abuse learned outside the confessional.[52] Washington's approach of abrogating the privilege for confessional communications is unusual, and the constitutional question is unresolved as of this writing.[53]

The Washington approach may be an outlier today, but the mandatory-reporting pressure it embodies is not, and that pressure carries a structural implication for AI-mediated spiritual communications that courts and commentators have not yet examined. As civil-law pressure on the confessional seal grows, some users will turn to religious AI platforms, believing them to offer more privacy. That belief is wrong, as AI-mediated spiritual communications currently have no privilege protection at all.[54] They

49. S.B. 5375, 69th Leg., Reg. Sess. (Wash. 2025).

50. Murphy, *supra* note 27 (reporting the DOJ Civil Rights Division investigation into Washington SB 5375); Emp. Div. v. Smith, 494 U.S. 872, 890 (1990) (establishing that neutral, generally applicable laws that incidentally burden religious practice do not violate the Free Exercise Clause); Fulton v. City of Phila., 593 U.S. 522, 533 (2021) (clarifying that laws that are not neutral or generally applicable are subject to strict scrutiny under the Free Exercise Clause).

51. Emp. Div. v. Smith, 494 U.S. 872, 878–79 (1990) (holding that the Free Exercise Clause does not exempt religious conduct from a neutral, generally applicable law); Fulton v. City of Phila., 593 U.S. 522, 533–34 (2021) (holding that a law is not generally applicable, and is thus subject to strict scrutiny, where it permits individualized, discretionary exemptions).

52. Dick, *supra* note 4, at 88–91 (surveying state mandatory reporting statutes and their standard approach of preserving the clergy privilege while requiring disclosure of abuse learned outside the confessional context).

53. Murphy, *supra* note 27.

54. CAL. EVID. CODE § 1030 (West 2026); N.Y. C.P.L.R. § 4505 (McKinney 2025) (no state clergy-penitent privilege statute encompasses an AI platform as a

footnote continued on next page

are not more protected than confessional communications under legislative attack. They are categorically less protected. Recognizing and closing this gap is a core purpose of the Digital Seal framework presented in this Article.

E. Applicability to Mistaken Identity

There does not appear to be any case law regarding whether the clergy-penitent privilege applies when a person mistakenly believes they are communicating with a priest. However, a similar situation regarding the attorney-client privilege may suggest that in such a case, the clergy-penitent privilege should apply. Courts have held that the attorney-client privilege applies even when the communication is between a person talking to someone he mistakenly believes is a lawyer.[55] In *United States v. Rivera*[56] the Southern District of New York explained that "It is common ground among the parties that the attorney-client privilege attaches to confidential communications made to an individual in the genuine, but mistaken, belief that he is an attorney."[57] Another court even maintained that the attorney-client privilege applies when an inmate reasonably believed that his cellmate was a lawyer.[58] Together, these cases reflect a putative professional standard: the privilege attaches based on the communicant's honest

qualifying privilege-holder); United States v. Heppner, 820 F. Supp. 3d 292, 297 (S.D.N.Y. Feb. 17, 2026) (holding directly that communications made to a commercial AI platform are not privileged regardless of content or the user's subjective expectation of confidentiality).

55. *See* Grace M. Giesel, *Upjohn Warnings, the Attorney-Client Privilege, and Principles of Lawyer Ethics: Achieving Harmony*, 65 U. Miami L. Rev. 109, 140 (2010) ("In addition to applying the honest-and-reasonable-belief analysis in the attorney-client privilege setting to the question of representational relationship, courts have also applied the analysis when the person consulted is not, in fact, a lawyer. The courts addressing this issue have stated that the privilege applies to a communication even if the person consulted is not admitted to any bar and has enjoyed no legal training. If the putative client honestly and reasonably believes that the person consulted is a lawyer, and if the other requirements of the privilege are satisfied, the privilege applies even though the person consulted is, in fact, not a lawyer.").

56. 837 F. Supp. 565 (S.D.N.Y. 1993).

57. *Id.* at 567 n.1.

58. United States v. Tyler, 745 F. Supp. 423, 425 (W.D. Mich. 1990).

and reasonable belief about who is receiving the confidence, not on the listener's actual credentials.

The putative professional standard as applied to attorney-client privilege provides a basis to argue that, in certain circumstances, a communication with Magisterium AI should be covered under the clergy-penitent privilege. Given that Magisterium AI markets itself as the "World's #1 answer engine for the Catholic Church," has a lack of any prominent disclaimer, and was trained exclusively on official Catholic documents, users could mistakenly believe that it is an official product of the Catholic Church. Some users who are technologically challenged may even falsely believe that a human member of the Catholic clergy is writing the responses. When asked, "Does the clergy-penitent privilege apply to this conversation?" it does not help that Magisterium AI answers, in part, that it "depends on secular law in your jurisdiction."[59]

The mistaken identity doctrine thus presents a colorable but ultimately insufficient basis for extending the clergy-penitent privilege to Magisterium AI communications. The attorney-client cases establish that courts will protect a communication when the privilege's underlying rationale—candor enabled by a reasonable expectation of confidentiality within a professional relationship—is present even if the specific credential is absent. That rationale applies with some force to a user who reasonably, if incorrectly, believes that Magisterium AI represents an authoritative voice of the institutional Catholic Church. It applies with less force, however, as the user's sophistication increases and the platform's commercial character becomes apparent. More fundamentally, the mistaken identity doctrine addresses only the first of the three independent doctrinal barriers that stand between AI-mediated religious communications and privilege protection. Even if a court were persuaded that a particular user held a genuine and reasonable belief that Magisterium AI functioned as a clerical advisor, the Third-Party Doctrine would still render those communications accessible to government process,

59. Casey Rockwell & Michael Conklin, Magisterium AI, "Does the clergy-penitent privilege apply to this conversation?" (May 21, 2026) (on file with author).

and *Heppner* would still foreclose the privilege claim as a matter of evidentiary law.

The mistaken identity argument is therefore best understood not as an independent path to privilege protection, but as evidence of the gap the Digital Seal is designed to close: users are already treating Magisterium AI as a spiritually authoritative confidant, and the law has no adequate response to that reality under existing doctrine. Taken together, the constitutional, canonical, and putative-professional foundations surveyed in this Part all point toward protecting the kind of disclosure Magisterium AI now receives, yet none of them is sufficient on its own to carry that protection across the line.

IV. Three Independent Doctrinal Barriers to Privilege Recognition

A. The Human-Advisor Requirement

Every U.S. state has a clergy-penitent privilege statute.[60] The privilege arises from the relationship, not the information.[61] Thus, the basic barrier here is merely definitional. An AI system is not a member of the clergy, a minister of religion, or a "similar functionary of a church or of a religious denomination" under any existing statute, no matter how sophisticated the AI becomes.[62] The barrier is therefore one of status rather than capability.

Some who seek to defend religious disclosures to religious chatbots reach to clergy-penitent privilege for guidance. The argument runs as an analogy because if the law shields what a penitent confides to a priest from compelled disclosure then it should likewise shield confidential disclosures to an AI "confessor" occupying the same or similar role. A 2025 article in the *Harvard Journal of Law and Technology* stated this precisely: the clergy-penitent analogy fails because it "presumes a religious vocation and shared sacramental understanding,"

60. Christine P. Bartholomew, *Exorcising the Clergy Privilege*, 103 Va. L. Rev. 1015, 1020–21 (2017).

61. *See* Dick, *supra* note 4, at 81 (discussing historical development of clergy-penitent privilege based on relationship); *see also* Aronson, *supra* note 37, at 591–93 (noting that all recognized privileges are relational in nature—they protect specific types of relationships, not specific categories of information).

62. Cal. Evid. Code § 1030 (West 2026); N.Y. C.P.L.R. § 4505 (McKinney 2025).

and "what is missing is the relationship with legally cognizable duties and institutional accountability." [63] Without those elements, "functional similarity is not a bridge to privilege; it is a mirage."[64]

This critique is correct under existing doctrine. It is also the most tractable of the three barriers discussed in this Part because it is a statutory issue, not a constitutional one. Congress or a state legislature can define a new privilege-holder category covering AI platforms that operate under defined institutional religious conditions. No constitutional amendment is needed. No precedent must be overruled. Parts V and VI of this Article develop that argument into the Digital Seal proposal.

However, one prior article reaches the opposite conclusion on this point and warrants direct engagement. Samuel Dick argues that existing statutes are already broad enough to cover AI religious communications without new legislation, pointing to language such as California Evidence Code section 1030's reference to "a similar functionary of a church or of a religious denomination or religious organization" and to cases where courts have construed "clergy" to encompass non-ordained spiritual counselors in functional roles. [65] Courts have sometimes read "clergy" by function rather than by title. In *In re Verplank*, [66] for instance, a federal court held that a non-ordained counselor staffing a minister-led counseling program fell within the clergy-penitent privilege because their work conformed in a general way to that already established by ministers and they acted

63. Ira P. Robbins, *Against an AI Privilege*, HARV. J.L. & TECH. DIG. (Nov. 7, 2025), https://jolt.law.harvard.edu/digest/against-an-ai-privilege [https://perma.cc/QY3M-A2NV].

64. *Id.*

65. *Compare* Dick, *supra* note 4, at 97–111 (arguing that the privilege could reach AI by pairing "similar functionary" language with judicial deference to a church's own designation of its clergy and the doctrine of First Amendment personhood) *with* CAL. EVID. CODE § 1030 (West 2026) (arguing that broad statutory language such as California Evidence Code § 1030's reference to "a similar functionary of a church or of a religious denomination or religious organization" is capacious enough to encompass AI platforms operating under formal ecclesiastical oversight, and citing cases in which courts have construed "clergy" expansively beyond ordained ministers of traditional denominations).

66. *In re* Verplank, 329 F. Supp. 433 (C.D. Cal. 1971).

under a minister's authority.[67] Under Samuel Dick's reading, an AI platform operating under formal ecclesiastical oversight would qualify as a "similar functionary" and its users' communications would already be protected.[68]

Nevertheless, Samuel Dick's statutory argument has surface appeal but fails on examination for two independent reasons. First, the cases he cites extending "clergy" to functional counselors all involved human beings exercising relational and pastoral judgment. The functional role that justified privilege in those cases was human counsel, not information retrieval. No reported decision has extended 'functionary' to a commercial software product, and the ordinary meaning of that term in context connotes a person, not a program.

Second, and more fundamentally, *Heppner* resolves the practical question Samuel Dick's argument leaves open. Even if a creative statutory argument could place an AI platform within the definition of "clergy," *Heppner* establishes that communications made to a commercial AI platform are not privileged as a matter of the evidentiary law governing those communications. The statutory definitional argument and the evidentiary privilege analysis are distinct, and *Heppner* bars the privilege regardless of how the statutory definition is construed. New legislation is not merely preferable to Samuel Dick's reading of existing law, it is necessary.

B. The Third-Party Doctrine

The second barrier is the Third-Party Doctrine. Information voluntarily shared with a third party, or stored on a third party's servers, loses Fourth Amendment protection and is generally

67. *Id.* at 436.

68. *See, e.g.*, *In re* Grand Jury Investigation, 918 F.2d 374, 377 (3d Cir. 1990) (construing the federal common-law clergy privilege to attach to a communication made to a lay associate in a religious community where the communication was made in the associate's capacity as a spiritual counselor); Kos v. State, 15 S.W.3d 633, 639 (Tex. App. 2000) (holding privilege does not apply because the communication was not made by the defendant seeking spiritual guidance).

accessible to government process without a warrant.[69] When a user submits a query to Magisterium AI, that query travels through Longbeard's servers. The company possesses the communication. Whatever contractual promises Longbeard makes about data handling, the Fourth Amendment does not shield that information from a government subpoena.[70] The privilege analysis is a distinct question from the Fourth Amendment analysis, but the structural problem is the same: a commercial intermediary stands between the user and any protection the law might otherwise provide.[71]

Courts have grappled with this structure in AI-related privacy litigation.[72] The key factual distinction courts draw is between data processed only to serve the communicating user and data used for the provider's own commercial purposes. AI platform terms of service routinely blur this line.[73] Regardless of how that line is drawn in a given case, the commercial company possesses the communications, and no current privilege shields them from legal process.

The scope of the Third-Party Doctrine in the digital context is actively contested. *Carpenter v. United States* narrowed the doctrine for long-term cell-site location data but expressly declined to disturb *Smith v. Maryland* and *U.S. v. Miller* or to articulate a general rule for

69. *E.g.*, Smith v. Maryland, 442 U.S. 735, 743–44 (1979); United States v. Miller, 425 U.S. 435, 443 (1976). The Third-Party Doctrine has been cabined somewhat by Carpenter v. United States, 585 U.S. 296, 309–10, 315–16 (2018), but the Court in *Carpenter* declined to disturb *Smith* and *Miller* or to establish a general rule protecting all third-party-held records.

70. *Miller*, 425 U.S. at 443–44.

71. HALLOW HELP & SUPPORT CTR., *supra* note 6 (stating that "user data is never used to train the model"—a contractual representation that may inform a court's equitable analysis but does not establish a Fourth Amendment protection or an evidentiary privilege).

72. Molly Jennings et. al, *Year in Review: 2025 Artificial Intelligence-Privacy Litigation Trends*, WILMERHALE (Feb. 2, 2026), https://www.wilmerhale.com/en/insights/blogs/wilmerhale-privacy-and-cybersecurity-law/20260202-year-in-review-2025-artificial-intelligence-privacy-litigation-trends [https://perma.cc/PB6G-BU29].

73. *Id.* (discussing the threshold distinction courts draw in CIPA AI-eavesdropping litigation between data processed exclusively to benefit the communicating consumer and data leveraged for the provider's independent commercial purposes).

third-party-held digital records.[74] As of May 2026, the Supreme Court is currently considering that question in *Chatrie v. United States*,[75] argued April 27, 2026, which asks whether a geofence warrant directing Google to search its users' location history constitutes a Fourth Amendment search.[76] The Fourth Circuit split 7-7 en banc on whether a search occurred, and the Supreme Court granted certiorari limited to that threshold question.[77]

The *Chatrie* oral argument carries two direct implications for religious AI communications. First, multiple Justices appeared to assume that data held by cloud-based services like email carries warrant protection even when voluntarily shared with a service provider, suggesting possible willingness to limit the doctrine's reach in digital contexts. [78] If that reasoning prevails, AI platform

74. *Carpenter*, 585 U.S. at 316 ("Our decision today is a narrow one. We do not express a view on matters not before us . . . We do not disturb the application of *Smith* and *Miller* or call into question conventional surveillance techniques and tools . . . Nor do we address other business records that might incidentally reveal location information.").

75. Chatrie v. United States, 146 S. Ct. 2193 (2026).

76. *Id.* at 2196. The Court granted certiorari on January 16, 2026, limited to Question 1: whether the execution of a geofence warrant directed at a third-party service provider constitutes a Fourth Amendment search. *See* Chatrie v. United States, SCOTUSBLOG, https://www.scotusblog.com/cases/chatrie-v-united-states/ [https://perma.cc/RNN8-3GNU] (last visited May 19, 2026).

77. *See* Mailyn Fidler, *Digital Location Data Heads Back to the Supreme Court*, SCOTUSBLOG (Apr. 24, 2026), https://www.scotusblog.com/2026/04/digital-location-data-heads-back-to-the-supreme-court-/ [https://perma.cc/4MR5-ANH7]; Ben Evelev & Olivia Parker, *Oral Argument Preview:* Chatrie v. United States, LAWFARE (Apr. 24, 2026, at 11:59 ET), https://www.lawfaremedia.org/article/oral-argument-preview-chatrie-v.-united-states [https://perma.cc/C8NW-UBNC].

78. *See* Jake Laperruque, *SCOTUS's Signals in* Chatrie *and on the Potential Limits of Location Tracking*, TECH POLICY PRESS (Apr. 29, 2026), https://www.techpolicy.press/scotuss-signals-in-chatrie-and-on-the-potential-limits-of-location-tracking/ [https://perma.cc/2HLY-Q6N2] (noting that "the Justices seemed to take for granted the notion that data held by common cloud-based services (like email and online calendars) were entitled to warrant protections despite being voluntarily shared with a third party," and that "[e]ven the government conceded that data users share via email and other third-party digital services likely couldn't be treated as voluntarily shared and freely grabbed by police.").

communications may eventually require a warrant to access. That would remove one of the three barriers identified in this Part, but not the other two. The human-advisor requirement and *Heppner* are entirely independent of the Fourth Amendment analysis, and the Digital Seal framework would remain necessary on those grounds alone.

Second, the Brennan Center's amicus brief warned that permitting reverse warrants without Fourth Amendment scrutiny would expose data collected near houses of worship and other sensitive religious locations.[79] Religious AI platforms present the same problem in digital form: a government subpoena to Magisterium AI reaches the same spiritual disclosures a physical search of a confessional would reach, but triggers none of the constitutional analysis a physical search requires. Whatever the Court decides in *Chatrie*, that structural asymmetry between physical and digital religious communications is the core problem the Digital Seal is designed to address.

The Third-Party Doctrine therefore operates on religious AI communications in two distinct ways. As a matter of Fourth Amendment law, its application to platform-held records is uncertain. As a matter of evidentiary law, however, that uncertainty offers users no present protection. Whatever the doctrine's eventual scope, a commercial intermediary holds the communication from the moment it is transmitted, and the absence of a constitutional barrier to compelled production is compounded by the absence of any privilege that would resist it.

79. *See* Brief of the Center for Democracy and Technology et al. as Amici Curiae Supporting Petitioner at 3, Chatrie v. United States, 146 S. Ct. 2193 (2026) (No. 25-112), https://www.supremecourt.gov/DocketPDF/25/25-112/399465/20260227091528269_2026-02-27-CDT_Amicus%20Brief_Chatrie_UnitedStates.pdf [https://perma.cc/6BWD-Z9Y3] ("Such a ruling permitting the use of geofence warrants could also open the door to the use of other novel reverse warrants that pose broad risks to individuals' privacy."). The intersection with religious AI communications is direct: a geofence warrant covering a house of worship collects location data; a subpoena to a religious AI platform collects content data. Both represent government access to sensitive religious information held by a third party without any established privilege protection.

C. *United States v. Heppner and the Foreclosure of AI Privilege Claims*

The third barrier is the most recent. In *United States v. Heppner*, a 2025 federal prosecution of a corporate executive for securities and wire fraud arising from his role at GWG Holdings, the district court faced a nationwide case of first impression.[80] The question was whether communications with generative AI platforms could be shielded from discovery. After his grand jury indictment, Heppner used a commercially available version of Claude to prepare documents in anticipation of trial.[81] The queries were fed by information that he had received from his lawyers and sharing the results of his queries with them.[82] The FBI seized those materials during a search.[83]

The court rejected the privilege claim. It held that "non-privileged communications are not somehow alchemically changed into privileged ones upon being shared with counsel," and that communications made to a commercial AI platform are not privileged regardless of content or the user's subjective expectations.[84] The court also declined to create any new AI-generated privilege category.[85]

This holding comes at a pivotal time. As early as 2023, studies revealed that large numbers of people were using generative AI models for legal advice.[86] Existing case law establishes that attorney-client privilege may be extended to include communications with a legal

80. United States v. Heppner, 820 F. Supp. 3d 292, 294–95 (S.D.N.Y. 2026); *see also* Zagros S. Bassirian & William D. Doodokyan, *Think Before You Chat: Using AI for Legal Questions May Waive Privilege and Create Discovery Risks*, NAT'L L. REV. (Apr. 15, 2026), https://natlawreview.com/article/think-you-chat-using-ai-legal-questions-may-waive-privilege-and-create-discovery [https://perma.cc/W2FX-3LMG] (discussing *Heppner*).

81. *Heppner*, 820 F. Supp. 3d at 295.

82. *Id.*

83. *Id.*

84. *Id.* at 296–97.

85. *See id.* at 296–97.

86. Steve Carroll, *New Survey Identifies How Consumers Would Be Willing to Use Generative AI to Address Legal Needs*, LEXISNEXIS (May 4, 2023), https://www.lexisnexis.com/community/insights/legal/b/thought-leadership/posts/new-survey-identifies-how-consumers-would-be-willing-to-use-generative-ai-to-address-legal-needs [https://perma.cc/K5PU-AARC].

assistant,[87] or an interpreter.[88] However, to be effective, it must ultimately be rooted in a human lawyer.[89] Some have recently argued that pro se litigants who ask generative AI models for legal advice should receive protections under a work-product theory, but existing case law does not recognize such a privilege.[90]

Heppner involved attorney-client privilege. Its holding is technically limited to that context. But its reasoning applies directly to any AI privilege claim: a commercial AI platform is not a recognized privilege-holder, a user's subjective expectation of privacy does not create privilege where none exists in law, and AI-generated communications held by a third-party company are discoverable. A court considering a clergy-penitent privilege claim over Magisterium AI communications would likely reach the same result under current doctrine.

V. THE SCHOLARLY DEBATE OVER AN AI INTERACTION PRIVILEGE

Legal scholars have begun addressing the gap between existing privilege doctrine and AI-mediated communications. Two camps have

87. *See* PAUL S. MILICH, GEORGIA RULES OF EVIDENCE § 21:3 (2024–2025 ed.) ("The scope of the attorney-client relationship includes, by necessity, the network of agents and employees of both the attorney and client, acting under the direction of their respective principals, to facilitate the legal representation. The attorney's privileged network includes administrative and legal assistants, from file clerks to law clerks.") (footnote omitted).

88. *See* Walsh v. CSG Partners, LLC, 544 F. Supp. 3d 389, 392 (S.D.N.Y. 2021) (explaining that the attorney-client privilege is not waived despite the presence of a third party "if the purpose of the third party's participation is to improve the comprehension of the communications between attorney and client" (quoting United States v. Ackert, 169 F.3d 136, 139 (2d Cir. 1999))).

89. *See* TEX. R. EVID. 503(a)(4)(A), (b)(1)(A) (permitting the attorney-client privilege for communications between the client "and the client's lawyer or the lawyer's representative" and defining a "lawyer's representative" as "one employed by the lawyer to assist in the rendition of "professional legal services").

90. Anoo Dinesh Vyas, *Confidentiality of AI Conversations: Protecting Self-Represented Litigants Who Use ChatGPT for Legal Advice*, 26 DUKE L. & TECH. REV. 186, 186 (2026) (supporting this position in part by emphasizing the effects on access to justice for litigants who cannot afford an attorney).

formed, and their disagreement shapes the design of the Digital Seal framework presented in this Article.

The first camp argues for a general AI interaction privilege built on the *Jaffee* model. Users now treat AI systems as confidants across medical, legal, and spiritual contexts. These advocates argue that the social benefit of candid disclosure in those settings outweighs the evidentiary cost of protection.[91] The *Jaffee* analogy is direct: the Supreme Court recognized the psychotherapist-patient privilege in 1996 because therapeutic candor produces social value.[92] Courts and legislatures should apply the same logic to AI communications serving comparable functions.[93]

The second camp, promoted by Ira P. Robbins, argues that AI lacks the two things privileges actually require: "professional role" and "institutional accountability."[94] This camp also notes that the Supreme Court has consistently refused to recognize new privileges in the absence of strong state-law consensus and clear social necessity, citing the Court's rejection of an academic peer review privilege in *University of Pennsylvania v. EEOC* as an example.[95] The University resisted an EEOC subpoena seeking confidential peer-review materials generated during the tenure process.[96] In defense, they argued that compelled disclosure would undermine the candor necessary for an effective

91. *See*, *e.g.*, Jennings, *supra* note 72 (surveying the range of AI privacy claims); *see also* Kirk J. Nahra, Molly Jennings & Ali A. Jessani, *Year in Review: 2024 Generative AI Litigation Trends*, WilmerHale (Mar. 7, 2025), https://www.wilmerhale.com/en/insights/blogs/wilmerhale-privacy-and-cybersecurity-law/20250307-year-in-review-2024-generative-ai-litigation-trends [https://perma.cc/H4UM-FGNN] (surveying 2024 generative AI privacy litigation).

92. Jaffee v. Redmond, 518 U.S. 1, 10–11 (1996).

93. *See* Melanie Stephens Stone, Case Note, Jaffee v. Redmond: *The Supreme Court Adopts a Federal Psychotherapist-Patient Privilege and Extends the Scope to Encompass Licensed Social Workers*, 48 Mercer L. Rev. 1283, 1290 (1997) (describing the *Jaffee* Court's reliance on the functional value of confidential therapeutic relationships as the basis for privilege recognition).

94. Robbins, *supra* note 63.

95. *Id.*; Univ. of Pa. v. EEOC, 493 U.S. 182, 189 (1990) (rejecting an asserted academic peer review privilege under Federal Rule of Evidence 501, notwithstanding the university's argument that peer review candor would be chilled by compelled disclosure).

96. *Univ. of Pa.*, 493 U.S. at 189–95.

process.[97] The Supreme Court declined to recognize a federal common-law privilege protecting those materials and went on to emphasize that courts should be cautious in creating new evidentiary privilege absent compelling justification.[98] Functional similarity to a human professional is not enough; the full institutional apparatus of licensure, fiduciary duty, and regulatory oversight must be present.

This Article accepts the second camp's diagnosis. A general AI interaction privilege covering all chatbot communications would be unworkable and overbroad. However, the second camp's reasoning loses force when applied to AI platforms that are structurally embedded within recognized religious institutions subject to canonical or equivalent accountability. The problem those critics identify, the absence of institutional accountability, is exactly what the Digital Seal's first structural condition addresses. The argument this Article makes is not about functional similarity; rather, it is about institutional structure, which is the criterion the critics correctly identify as decisive.

The relevant test remains *Jaffee*. Federal Rule of Evidence 501 permits courts to recognize new privileges based on "reason and experience."[99] *Jaffee* required near-uniform state-law recognition and a demonstrated social good dependent on confidential communication.[100] The Digital Seal framework does not presently satisfy both prongs. Near-uniform state recognition does not yet exist because no state has enacted a religious AI privilege. This is not a reason to abandon the *Jaffee* framework. It is a reason to pursue the legislative path rather than the judicial one. Federal Rule of Evidence 501 authorizes judicial recognition of new privileges, but Congress and state legislatures may create privileges independently of Rule 501 and

97. *Id.*

98. *Id.*

99. FED. R. EVID. 501.

100. Jaffee v. Redmond, 518 U.S. 1, 10–13 (1996) (relying on near-uniform state recognition and a demonstrated social good dependent on confidential communication in recognizing a new privilege under Federal Rule of Evidence 501); Amann & Imwinkelried, *supra* note 3, at 1025–28.

without satisfying *Jaffee's* consensus requirement. [101] The psychotherapist-patient privilege existed in state legislative consensus before *Jaffee*. The Court in *Jaffee* recognized it judicially because the legislative and professional consensus was already overwhelming.[102] Judicial recognition marked the end of that privilege's development, not its beginning.

The Digital Seal proposal posited in this Article reverses that sequence, seeking legislative action to build the consensus that judicial recognition would eventually require. Robbins' critique of an AI privilege recognized by courts does not reach a privilege created by statute.[103] The critique is that courts should not invent new privilege categories without the institutional accountability structures that existing privileges require. [104] Statutes can supply those structures directly, and that is precisely what the Digital Seal's three conditions do.

---

101. FED. R. EVID. 501 (providing that the common law governs privilege claims "unless any of the following provides otherwise: the United States Constitution; a federal statute; or rules prescribed by the Supreme Court"). Congress's authority to create statutory privileges independent of the *Jaffee* common-law consensus framework is confirmed by Rule 501's legislative history. *See* S. REP. NO. 93-1277, at 13 (1974), *as reprinted in* 1974 U.S.C.C.A.N. 7051, 7059 (stating that Congress's adoption of Federal Rule of Evidence 501 "should not be understood as disapproving any recognition" of specific privileges enumerated in the original proposed rules, and that "the recognition of a privilege based on a confidential relationship and other privileges should be determined on a case-by-case basis"); *see also* Amann & Imwinkelried, *supra* note 3, at 1026–27 (noting that the psychotherapist-patient privilege existed in state regulatory consensus before *Jaffee* formally recognized it under Federal Rule of Evidence 501, and that the Court treated that state legislative experience as a basis for federal recognition).

102. *Jaffee*, 518 U.S. at 12–13 ("[A]ll 50 States and the District of Columbia have enacted into law some form of psychotherapist privilege," and "[i]t is of no consequence that recognition of the privilege in the vast majority of States is the product of legislative action rather than judicial decision."); *see also* Amann & Imwinkelried, *supra* note 3, at 1025–28 (analyzing how the *Jaffee* Court treated prior state legislative action as the evidentiary basis for Rule 501 recognition).

103. Robbins, *supra* note 63 ("Legislatures may experiment with narrow protections for therapeutic uses, but courts should not expand the common law of privilege.").

104. *Id.*

VI. The Digital Seal: A Proposed Statutory Framework

A. Design Principles

The Digital Seal is a proposed qualified statutory evidentiary privilege for AI-mediated religious communications satisfying three structural conditions. The Digital Seal framework is governed by four design principles.

First, the framework is confessionally neutral. It does not favor any religious tradition or treat Catholic canonical practice as the template. Any tradition whose AI platform satisfies the structural criteria qualifies. Each institution decides internally what forms of AI-mediated spiritual communication, if any, it wants to bring within the framework.[105]

Second, the privilege is institution-dependent, not function-dependent. What matters is what the platform is, namely an AI system formally anchored within a recognized religious institution, not what the communication contains or what role the AI plays in a particular user's spiritual life. This directly responds to the second camp's institutional accountability objection.[106]

Third, the privilege is qualified, not absolute. This Article does not propose replicating the absolute character of the canonical *sigillum sacramentale* in civil law. That would be neither constitutionally required nor wise. The privilege is qualified and subject to defined exceptions.[107]

Fourth, the framework is neutral. It takes no position on whether AI spiritual counsel is theologically appropriate. Instead, it creates legal conditions under which such counsel can be received without the risk of compelled disclosure. This neutrality is practical, as it treats all

105. *See* Schragger, Schwartzman & Tebbe, *supra* note 35, at 209 (describing the emerging rule that government support of religion is permitted "so long as it is neutral, as a formal matter, between religion and nonreligion and among religions").

106. Robbins, *supra* note 63.

107. *Cf. Jaffee*, 518 U.S. at 17–18 & n.19 (rejecting a case-by-case balancing component that "would eviscerate the effectiveness of the privilege," while declining to "delineate its full contours" and noting that "there are situations in which the privilege must give way").

religious traditions equally and does not venture into the highly contested domain of theological legitimacy. This restraint is consistent with Supreme Court precedent that generally does not consider whether a sincerely held religious conviction is the correct interpretation of a given religion's orthodoxy.[108]

These four principles are interdependent rather than merely cumulative. Confessional and theological neutrality keep the framework within the competence of civil law and ensure that legislatures resolve questions of institutional structure rather than questions of religious truth. The institution-dependent criterion supplies the accountability that privilege law has always demanded, and the privilege's qualified character keeps that accountability enforceable by preserving the exceptions that make evidentiary privilege administrable.

B. The Three Structural Conditions

The Digital Seal privilege applies to a communication made to an AI platform that meets all of the following three conditions at the time of the communication. The first condition is institutional religious anchoring. The platform must be formally organized under, or formally affiliated with, a recognized religious institution. That institution must exercise ongoing doctrinal oversight over the platform's outputs and must have formally adopted the platform as an authorized resource.[109] Formal affiliation requires a documented,

108. *See*, *e.g.*, Burwell v. Hobby Lobby Stores, Inc., 573 U.S. 682, 725 (2014) ("[I]t is not for us to say that their religious beliefs are mistaken or insubstantial."); Hernandez v. Comm'r, 490 U.S. 680, 699 (1989) ("It is not within the judicial ken to question the centrality of particular beliefs or practices to a faith, or the validity of particular litigants' interpretations of those creeds."); Thomas v. Rev. Bd. of the Ind. Emp. Sec. Div., 450 U.S. 707, 716 (1981) ("Courts are not arbiters of scriptural interpretation."); United States v. Ballard, 322 U.S. 78, 86 (1944) ("Heresy trials are foreign to our Constitution. Men may believe what they cannot prove. They may not be put to the proof of their religious doctrines or beliefs.").

109. The term "recognized religious institution" is intended to track existing statutory language. *See*, *e.g.*, Cal. Evid. Code § 1030 (West 2024) (defining clergy privilege-holder to include "a similar functionary of a church or of a religious denomination or religious organization"); N.Y. C.P.L.R. § 4505 (McKinney 2025).

publicly disclosed, bilateral commitment: the platform agrees to operate within the institution's doctrinal parameters, and the institution commits to ongoing oversight. An advisory committee without binding authority is not sufficient. In its current form, Magisterium AI's scholarly advisory committee structure likely does not satisfy this condition.[110] That committee was convened by Longbeard rather than by any ecclesiastical authority, and its purview extends to the curated corpus rather than to the responses users actually receive.[111] This condition does the same work in the Digital Seal framework that professional licensure and fiduciary duty do in the attorney-client and psychotherapist-patient contexts: it creates an institutional accountability structure that makes the privilege-holder legally trustworthy.

The second condition is data-minimization architecture. The platform must implement end-to-end encryption of all user communications during transmission and storage, a binding prohibition on use of those communications for model training or commercial analytics, a data retention limit of no more than thirty days, and an annual independent audit verifying compliance.[112] These requirements draw on the EU Artificial Intelligence Act's data-minimization principles and parallel requirements in U.S. state AI consumer protection statutes.[113] These requirements reflect a broader regulatory trend towards limiting AI systems' use of personal data to information that is necessary for a defined purpose. The EU AI Act

110. Ang, *supra* note 5, at 241 (describing the scholarly advisory committee structure). This Article expresses no view on whether Magisterium AI's existing institutional relationships are sufficient for any purpose other than the statutory privilege analysis proposed herein.

111. *Id.*

112. These requirements are modeled in part on the data-minimization principles of Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 on Artificial Intelligence (Artificial Intelligence Act); *see also AI Chatbots at the Crossroads: Navigating New Laws and Compliance Risks*, COOLEY LLP (Oct. 21, 2025), https://www.cooley.com/news/insight/2025/2025-10-21-ai-chatbots-at-the-crossroads-navigating-new-laws-and-compliance-risks [https://perma.cc/9WC5-6R95] (surveying state AI disclosure and data-minimization statutes including California's 2024 AI healthcare disclosure law, Colorado's CAIA, and Utah's AIPA).

113. *Id.*; Jennings, *supra* note 72.

reinforces this principle through data-governance requirements for high-risk systems, while emerging US state AI laws similarly emphasize transparency, data governance, and consumer control over information used in automated decision-making.[114] Their purpose is to reduce the commercial-intermediary problem. They do not resolve it entirely, but they limit the practical risk that user communications will remain accessible after compelled disclosure is ordered.[115]

The third condition is ecclesiastical accountability. The affiliated religious institution must designate a named individual with canonical or equivalent institutional authority who is responsible for the platform's doctrinal accuracy and is available to address user complaints.[116] Antiqua et Nova states that "only the human person can be morally responsible."[117] This condition gives that principle legal form. It is not a theological requirement imposed by civil law. It is a structural prerequisite grounded in the same institutional-accountability logic that underlies every existing privilege.

The three conditions are cumulative rather than alternative. A platform satisfying only the first is a doctrinally supervised commercial product whose user data remains fully discoverable. A platform satisfying only the second is a privacy-protective chatbot with no institutional warrant for what it teaches. And a platform satisfying only the third has designated a responsible party without giving them anything meaningful to oversee. Only in combination do the conditions reproduce the features that make existing privileges administrable.

---

114. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024, art. 10(2)(b), (e), 2024 O.J. (L 1689) 1. (EU) (requiring data-governance practices for high-risk AI systems, including attention to data collection, data origin, and the original purpose for which personal data were collected); Colo. Rev. Stat. § 6-1-1701 et seq. (as amended by S.B. 26-189, 75th Gen. Assemb., 2d Reg. Sess. (Colo. 2026)) (establishing disclosure, data-access, correction, and human-review protections for consumers affected by automated decision-making technologies); Jennings, *supra* note 72.

115. *Cf.* Jaffee v. Redmond, 518 U.S. 1, 17 (1996) ("A promise of confidentiality must be absolute to encourage the patient to be forthcoming.").

116. *Antiqua et Nova*, *supra* note 8, at ¶ 111 ("Only the human person can be morally responsible . . . .").

117. *Id.*

C. Scope, Exceptions, and the Mandatory Reporting Problem

The Digital Seal privilege has three exceptions. The first is the crime-fraud exception, modeled on its attorney-client privilege analog. Communications involving the planning or commission of a future crime or fraud are not protected. Courts may conduct in camera review to determine whether the exception applies.[118]

The second exception covers mandatory reporting obligations. Where a state has enacted a mandatory reporting statute requiring disclosure of child abuse or neglect information, such as Washington SB 5375, the Digital Seal privilege yields to that statute.[119] This applies the same policy judgment to AI-mediated and human-mediated spiritual communications equally.

The third exception covers imminent threats to identifiable persons. The Digital Seal privilege does not protect communications revealing an imminent, serious threat to a specific third party.[120] The Supreme Court left open a similar exception for the psychotherapist-patient privilege in *Jaffee*.[121]

A practical problem exists: an AI platform cannot exercise the moral judgment a human religious leader can. It cannot recognize on its own when a communication triggers one of these exceptions. The Digital Seal framework requires qualifying platforms to establish a human review protocol for communications that may implicate an exception. That human-in-the-loop requirement is consistent with Antiqua et Nova, the Vatican's 2025 doctrinal note on the relationship between artificial intelligence and human intelligence, which emphasizes that the human must remain responsible over AI outputs

118. *See* United States v. Zolin, 491 U.S. 554, 563 (1989) (discussing the crime-fraud exception to attorney-client privilege and the procedure for in camera review of challenged documents).

119. S.B. 5375, 69th Leg., Reg. Sess. (Wash. 2025).

120. *Id.*

121. *See Jaffee*, 518 U.S. at 18, n.19 (leaving open "the question of whether our recognition of the privilege . . . bars disclosure in cases . . . in which the patient threatens violence to a readily identifiable third person"). The Court's reservation suggests that even absolute-seeming privileges may yield to imminent-threat circumstances.

and provides a workable mechanism for reconciling the privilege with mandatory reporting laws.[122]

One structural objection to the human review protocol warrants direct response. If a qualifying platform must designate human reviewers to identify exception-triggering communications, a government seeking access to privileged communications might simply subpoena those reviewers rather than the platform itself, arguing that the privilege never attached or was waived by disclosure to a third party. That argument fails for the same reason it fails in the attorney-client context. The privilege attaches to the communication at the moment it is made, not to the medium or the person through whom it later passes. An attorney who reviews a client's documents does not destroy the privilege in those documents, and a priest who discloses a confession under compulsion does not retroactively strip the communication of its privileged character.[123]

The human reviewer in the Digital Seal framework occupies an analogous position: a designated agent of the privileged institution whose access to the communication is for the limited purpose of exception identification, not an independent third party whose involvement breaks the chain of confidentiality. Courts have recognized this principle in the attorney-client context through the *Kovel* doctrine, which extends privilege to third parties whose involvement facilitates rather than defeats the protected relationship.[124]

---

122. *Antiqua et Nova*, *supra* note 8, at ¶ 112 (insisting on human oversight and moral responsibility over AI outputs, particularly in sensitive domains).

123. *Zolin*, 491 U.S. at 563 (1989) (recognizing that the crime-fraud exception requires in camera review precisely because privilege presumptively attaches to the document itself, not to the hands currently holding it); *see also* 1983 CODE c.983, § 1 (canonical seal attaches to the communication irrespective of subsequent disclosure).

124. United States v. Kovel, 296 F.2d 918, 921–22 (2d Cir. 1961) (holding that attorney-client privilege extends to communications made to an accountant employed by a law firm where "the accountant is necessary, or at least highly useful, for the effective consultation between the client and the lawyer," and that such third-party involvement "should be no more than [the lawyer's] alter ego in these respects").

The Digital Seal framework codifies the same logic: the human review protocol is a structural requirement of the privilege, not a waiver of it, and a subpoena directed at the reviewer reaches no further than a subpoena directed at the platform. Human review is therefore better understood as a part of the privileged relationship than as an intrusion upon it.

D. Canonical Alignment, Constitutional Considerations, and the Limits of Civil Law

The Digital Seal framework does not conflict with canonical obligations. For the Catholic tradition, this Article does not propose a civil-law analog to the sacramental seal. That seal is absolute, theologically grounded, and admits no exceptions under canon law.[125] The Digital Seal is a qualified civil-law privilege for AI-mediated spiritual communications. For other religious traditions the framework is similarly complementary: it creates a legal structure that accommodates each tradition's internal standards without imposing a single doctrinal model. The civil privilege therefore operates as a floor rather than a ceiling; it protects what secular law is competent to protect and leaves each tradition free to impose stricter obligations.

The canonical foundation for the Digital Seal framework was significantly strengthened on May 25, 2026, when Pope Leo XIV issued *Magnifica Humanitas*, his encyclical on safeguarding the human person in the time of artificial intelligence.[126] The encyclical, the first full papal social encyclical to address AI directly, expressly calls for "adequate regulatory tools capable of upholding justice and curbing the distorting effects of technological power."[127] It warns that the

125. Apostolic Penitentiary Note, *supra* note 47. The absolute character of the canonical seal rests on theological foundations that are neither replicable nor appropriately addressed by civil legislation. This Article proposes a qualified civil-law privilege for a related but broader category of communications and makes no claim that such a privilege satisfies canonical obligations or replicates the *sigillum*.

126. Pope Leo XIV, *Magnifica Humanitas, Encyclical Letter on Safeguarding the Human Person in the Time of Artificial Intelligence* (May 25, 2026) [hereinafter Magnifica Humanitas], https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html [https://perma.cc/8JSE-SFTD].

127. *Id.* at ¶ 5 ("It is necessary to establish adequate regulatory tools capable of upholding justice and curbing the distorting effects of technological power.").

"main drivers of development are private, often transnational, parties" whose resources surpass those of many governments, making it "even more challenging to discern, govern and direct such power toward the common good."[128] The encyclical further observes that technology "takes on the characteristics of those who devise, finance, regulate and use it,"[129] and that the risk of dehumanization includes the reduction of "the mystery of the person into data and performance."[130] These are not merely theological propositions. They are the Catholic Church's own institutional judgment that AI operating in sensitive human domains requires regulatory governance that current U.S. law does not provide for religious AI communications.

For purposes of the Digital Seal framework advocated for in this Article, *Magnifica Humanitas* provides canonical authority for two of the framework's three structural conditions. The encyclical's call for regulatory accountability of private AI actors grounds the institutional anchoring condition: the Church itself demands that AI operating within its domain be subject to human oversight answerable to the common good, not merely to the commercial interests of its developer.[131] Its warning against reducing the person to "data and performance" grounds the data-minimization condition: treating a user's spiritual disclosures as training data or commercial analytics is

128. *Id.* at ¶ 5 ("Today, however, the main drivers of development are private, often transnational, parties that are endowed with resources and the capacity to intervene that surpass those of many Governments," making it "even more challenging to discern, govern and direct such power toward the common good.").

129. *Id.* at ¶ 9 ("In practice, however, technology is never neutral, because it takes on the characteristics of those who devise, finance, regulate and use it.").

130. *Id.* at ¶ 10 ("[T]he pretense that a single language—even a digital one—can translate everything, including the mystery of the person, into data and performance.").

131. *Id.* at ¶ 5 ("It is necessary to establish adequate regulatory tools capable of upholding justice and curbing the distorting effects of technological power" and noting that private transnational AI actors now hold "resources and the capacity to intervene that surpass those of many Governments," making it challenging "to discern, govern and direct such power toward the common good"); *id.* at ¶ 9 ("[T]echnology is never neutral, because it takes on the characteristics of those who devise, finance, regulate and use it.").

precisely what the encyclical identifies as dehumanizing.[132] The Digital Seal framework translates both of these canonical imperatives into the statutory form that U.S. courts and legislatures can act upon. Civil law cannot enforce a papal encyclical. But it can create a privilege framework whose structural requirements reflect the same institutional commitments the Catholic Church has articulated as necessary for AI to serve, rather than threaten, human dignity.

One constitutional objection deserves a direct response. Requiring formal affiliation with a recognized religious institution could be read as favoring organized religion over informal or individual religious practice. This raises Establishment Clause concerns.[133] However, this Article does not dismiss that concern. But the institutional anchoring requirement serves a secular and neutral purpose, ensuring the accountability structure that makes a privilege legally credible, and it applies across all religious traditions without singling out any subset for favorable treatment. That neutrality is consistent with the Supreme Court's recent First Amendment jurisprudence, which confirms that legislative accommodations of organized religious practice are permissible where the criteria are neutral and generally applicable.[134]

E. Comparative Law Grounding and Model Statutory Text

The Digital Seal framework does not arise in a comparative vacuum. Other jurisdictions have already recognized that religious belief data demands heightened legal protection, providing both

132. *Id.* at ¶ 10 ("[T]he pretense that a single language even a digital one—can translate everything, including the mystery of the person, into data and performance" constitutes the "Babel syndrome," and "[t]he risk of dehumanization—of building a future that excludes God and reduces the other to a means—is an ancient and ever-new temptation that today takes on a technical guise.").

133. *See* U.S. CONST. amend. I ("Congress shall make no law respecting an establishment of religion . . . .").

134. Schragger, Schwartzman & Tebbe, *supra* note 35, at 208–11 (analyzing the doctrine of "structural preferentialism" that permits legislative accommodations of organized religious practice where they rest on neutral, generally applicable criteria rather than denominational preference); *see also* Fulton v. City of Philadelphia, 593 U.S. 522, 541 (2021) (applying strict scrutiny and finding it unsatisfied because the City's contract allowed discretionary, individualized exemptions and was therefore not generally applicable).

precedent for the framework's data-minimization conditions and context for assessing the gap in U.S. law.

The European Union's General Data Protection Regulation classifies "personal data revealing religious or philosophical beliefs" as a special category under Article 9, subject to a general prohibition on processing unless a specific statutory exception applies.[135] The Court of Justice of the European Union extended this protection in 2022 to include data from which religious belief can be inferred, even where the belief is not explicitly stated. A religious AI chatbot operating in the EU that processes user queries touching on sacramental practice, moral theology, or doctrinal interpretation is already subject to heightened regulatory scrutiny that has no equivalent in United States law. This asymmetry is particularly significant because Magisterium AI operates globally. The same user query receives GDPR Article 9 protection in Frankfurt, Germany and no statutory protection in Fort Worth, Texas. The Digital Seal framework would begin to close that gap.

Canada presents a different but equally instructive case. Bill C-27, which would have created Canada's first comprehensive AI privacy statute and applied to Longbeard, died on the Order Paper in January 2025 when Parliament was prorogued.[136] Canada thus regulates

135. General Data Protection Regulation art. 9(1), Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 Apr. 2016 [GDPR] (prohibiting the processing of "personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs" unless a specific exception applies). In Case C-184/20, Vyriausioji tarnybinės etikos komisija, ECLI:EU:C:2022:601 (Aug. 1, 2022), the Court of Justice of the European Union held that the special-category prohibition extends to data from which religious belief can be inferred, not merely to data that explicitly identifies it. The European Commission has proposed amendments to Article 9 in the Digital Omnibus package (Nov. 2025) that would modify the inference-based rule, but those proposals remain under negotiation. *See European Commission Proposes Significant Reforms to GDPR, AI Act*, IAPP (Nov. 19, 2025).

136. Personal Information Protection and Electronic Documents Act, S.C. 2000, c. 5 (Can.) [PIPEDA]. Bill C-27, the Digital Charter Implementation Act, 2022, which would have replaced PIPEDA and introduced Canada's Artificial Intelligence and Data Act, died on the Order Paper in January 2025 when Parliament was prorogued. *See* Office of the Privacy Commissioner of Canada, *Prioritizing Privacy in a Data-Driven World: Annual Report 2024-*

footnote continued on next page

Magisterium AI under the Personal Information Protection and Electronic Documents Act ("PIPEDA"), a statute enacted in 2000 that predates both large language models and institutionally anchored religious AI.[137] PIPEDA's general consent requirements apply but provide no religious-belief special category. Quebec's Law 25, fully in force since September 2024, imposes stricter provincial requirements including mandatory privacy impact assessments for automated decision systems, but it also lacks a religious-belief designation. The result is that the company developing Magisterium AI operates in a jurisdiction that offers its users less statutory protection for religious communications than the EU does, and that the United States offers even less than Canada. Both the EU and Canadian experiences demonstrate that legislative action on this gap is achievable and that the institutional-anchoring model the Digital Seal proposes is consistent with how other jurisdictions approach sensitive data protection in AI contexts.

To provide a concrete starting point for legislative drafting, this Article proposes the following model statutory text for a Digital Seal provision, which could be enacted as a standalone federal statute or as an amendment to Federal Rule of Evidence 501:

> A communication made by a person to an institutionally anchored religious AI platform in connection with spiritual guidance, moral inquiry, or religious education shall be privileged and not subject

2025 (2025), https://www.priv.gc.ca/en/opc-actions-and-decisions/ar_index/202425/ar_202425/ [https://perma.cc/EC3A-DX9M] (noting Bill C-27's lapse). Canada thus operates Magisterium AI's home jurisdiction without a federal AI-specific privacy statute, relying on PIPEDA's general requirement that consent be obtained for collection of personal information. PIPEDA does not designate religious belief data as a special category requiring heightened protection, in contrast to the EU framework. Quebec's Act Respecting the Protection of Personal Information in the Private Sector (Law 25), fully in force since September 2024, imposes stricter provincial requirements including mandatory privacy impact assessments for AI-based automated decision systems, though it similarly lacks a religious-belief special category.

137. Personal Information Protection and Electronic Documents Act, S.C. 2000, c. 5, § 4(1)(a) (Can.) (applying Part 1 to organizations with respect to personal information they either collect, use, or disclose in the course of commercial activities).

> to compelled disclosure in any federal judicial or administrative proceeding. For purposes of this section, an institutionally anchored religious AI platform is an artificial intelligence system that (1) is formally authorized in writing by a recognized religious institution as an official resource of that institution; (2) is operated under the ongoing doctrinal oversight of a named individual with canonical or equivalent ecclesiastical authority designated by that institution; (3) implements end-to-end encryption, a prohibition on commercial use of user communications, and a data retention limit not to exceed thirty days; and (4) undergoes annual independent audit of compliance with the foregoing requirements. This privilege shall not apply to communications that constitute planning of a future crime or fraud, to information that a reporting statute requires to be disclosed, or to communications revealing an imminent threat of serious harm to an identifiable person.

This text is offered as a starting point, not a final proposal. The precise contours of each condition—the definition of "recognized religious institution," the scope of the mandatory reporting exception, and the audit mechanism—would require legislative refinement. The core structure, however, directly addresses each of the doctrinal barriers identified in Part IV. The remaining question is not whether such a privilege can be drafted, but whether it will be drafted before the first subpoena forces courts to answer the question without it.

VII. Conclusion

U.S. evidentiary law has expanded privileges when new social conditions required it. The psychotherapist-patient privilege did not exist in federal courts until 1996. The attorney-client privilege took decades of litigation to extend to in-house corporate counsel and has been applied to protect communications made to people mistakenly believed to be a lawyer. Each expansion reflected a recognition that a

new form of relationship was generating confidential communications that the existing framework did not adequately protect.

Institutional religious AI chatbots present the same kind of problem. This Article takes no position on the theological propriety of AI-mediated spiritual counsel. However, millions of users in more than 165 countries are using platforms like Magisterium AI, and many of them share personal moral and spiritual information in the process. U.S. law currently offers those users no protections. Under the Third-Party Doctrine, their communications belong to the platform. Under *Heppner*, those communications are discoverable. Under every state clergy-penitent privilege statute the absence of a human spiritual advisor ends the analysis. Three independent doctrines thus converge on a single result, and neither the sincerity of the user's expectation nor the sophistication of the platform's design alters it.

The novel Digital Seal framework proposed in this Article is a legal response to this gap. Rather than creating a sweeping privilege for all chatbot interactions, the framework ties protection to institutional accountability, data-minimization safeguards, and formal ecclesiastical oversight. In doing so, it attempts to preserve the core logic underlying traditional privileges while adapting that logic to emerging technological conditions. The proposal is intentionally limited, qualified, and religion-neutral.

Antiqua et Nova states that AI must remain a tool guided by human intelligence, oriented toward human dignity and the common good.[138] The Digital Seal framework provided in this Article translates that principle into statute. Whether legislatures ultimately adopt the Digital Seal or some alternative framework, the underlying issue will not disappear. As AI systems become increasingly integrated into religious life, the gap between users' expectations of confidentiality and the law's actual protections will continue to widen. The question is not whether courts and legislatures will eventually confront AI-mediated spiritual communications, but whether they will do so deliberately before a user's most private spiritual disclosure becomes evidence in court.

138. *Antiqua et Nova*, *supra* note 8, at ¶¶ 42, 112.